



**NAVAL
POSTGRADUATE
SCHOOL**

MONTEREY, CALIFORNIA

DISSERTATION

**A DATA-DRIVEN FRAMEWORK FOR RAPID MODELING
OF WIRELESS COMMUNICATION CHANNELS**

by

Douglas Horner

December 2013

Dissertation Supervisor:

Geoffrey Xie

Approved for public release; distribution is unlimited

THIS PAGE INTENTIONALLY LEFT BLANK

REPORT DOCUMENTATION PAGE			Form Approved OMB No. 0704-0188	
Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.				
1. AGENCY USE ONLY (Leave Blank)		2. REPORT DATE 12-20-2013		3. REPORT TYPE AND DATES COVERED Dissertation 06-30-2007 to 12-20-2013
4. TITLE AND SUBTITLE A DATA-DRIVEN FRAMEWORK FOR RAPID MODELING OF WIRELESS COMMUNICATION CHANNELS			5. FUNDING NUMBERS	
6. AUTHOR(S) Douglas Horner				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943			8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Department of the Navy			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this document are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government. IRB Protocol Number: N/A.				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited			12b. DISTRIBUTION CODE	
13. ABSTRACT (maximum 200 words) Accurate estimation and prediction of wireless signal strength holds the promise to improve a wide variety of applications in networking and unmanned systems. Current estimation approaches use either simplistic attenuation equations or detailed physical models that provide limited accuracy and may require a lengthy period of environmental assessment and computation. This dissertation presents a new, data-driven, stochastic framework for rapidly building accurate wireless connectivity maps. The framework advances the state of the art in three aspects. First, it augments the classic spatial interpolation procedure known as Kriging with a complementary additive approach to capture the typical anisotropic nature of wireless channels in cluttered environments. Second, it includes a technique for rapidly creating and maintaining a connectivity map in near real-time through the use of a spatial Bayesian recursive filter. Third, it introduces a novel methodology to adapt the resolution of a connectivity map based on the spatial characteristics and the quantity of available sample measurements. Detailed analyses, using several datasets collected recently in the Monterey Harbor, have confirmed the power and agility of the proposed approach.				
14. SUBJECT TERM Wireless Connectivity Maps, Random Fields, Kriging, Gaussian Process Models, ℓ_1 Regularized Logistic Regression, Kalman Filtering, Underwater Acoustic Networking			15. NUMBER OF PAGES 207	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFICATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UU	

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18

THIS PAGE INTENTIONALLY LEFT BLANK

Approved for public release; distribution is unlimited

**A DATA-DRIVEN FRAMEWORK FOR RAPID MODELING OF WIRELESS
COMMUNICATION CHANNELS**

Douglas Horner

B.S. and B.A., Boston University, 1985

M.S. and M.A., Naval Postgraduate School, 1997

Submitted in partial fulfillment of the
requirements for the degree of

DOCTOR OF PHILOSOPHY IN COMPUTER SCIENCE

from the

NAVAL POSTGRADUATE SCHOOL

December 2013

Author: Douglas Horner

Approved by: Geoffrey Xie
Professor of Computer Science
Dissertation Supervisor

Anthony Healey
Emeritus Distinguished Professor of
Mechanical and Aerospace
Engineering
Dissertation Committee Chair

Mathias Kolsch
Associate Professor of Computer
Science

Joel Young
Assistant Professor of Computer
Science

Timothy Chung
Assistant Professor of System
Engineering

John J. Leonard
Professor of Mechanical and Ocean
Engineering

Approved by: Peter Denning
Chair, Computer Science

Approved by: Douglas Moses
Associate Provost for Academic Affairs

THIS PAGE INTENTIONALLY LEFT BLANK

ABSTRACT

Accurate estimation and prediction of wireless signal strength holds the promise to improve a wide variety of applications in networking and unmanned systems. Current estimation approaches use either simplistic attenuation equations or detailed physical models that provide limited accuracy and may require a lengthy period of environmental assessment and computation. This dissertation presents a new, data-driven, stochastic framework for rapidly building accurate wireless connectivity maps. The framework advances the state of the art in three aspects. First, it augments the classic spatial interpolation procedure known as Kriging with a complementary additive approach to capture the typical anisotropic nature of wireless channels in cluttered environments. Second, it includes a technique for rapidly creating and maintaining a connectivity map in near real-time through the use of a spatial Bayesian recursive filter. Third, it introduces a novel methodology to adapt the resolution of a connectivity map based on the spatial characteristics and the quantity of available sample measurements. Detailed analyses, using several datasets collected recently in the Monterey Harbor, have confirmed the power and agility of the proposed approach.

THIS PAGE INTENTIONALLY LEFT BLANK

Table of Contents

1	Introduction	1
1.1	Network Control System	2
1.2	Distributed Coordination of a Wireless NCS	3
1.3	Feedback Control of Cooperative Systems	4
1.4	Communication Channel Estimation.	5
1.5	Main Hypothesis	6
1.6	Contribution of Dissertation	7
1.7	Organization of Dissertation	7
2	Background	9
2.1	Local Connectivity Map	9
2.2	Global Connectivity Map	12
2.3	Structure Learning with Markov and Conditional Random Fields	13
2.4	Cooperative Navigation.	15
2.5	Undersea Acoustic Networking.	17
3	The Local Connectivity Map	19
4	Structural Analysis	21
4.1	Stationarity	22
4.2	Structural Analysis	24
4.3	Conclusion.	33
5	Communication Considerations for Variogram Estimation	35
5.1	Simulated Semi-Variograms	35
5.2	Range Semi-Variograms	36
5.3	Curve Fitting the Empirical Range Semi-Variogram	37
5.4	Radial Semi-Variograms	38

5.5	Examples of Radial Semi-Variograms	41
6	Communications Kriging	47
6.1	Kriging	47
6.2	LCM Definition	62
6.3	Considerations for Communications	62
6.4	Undersea Acoustic Communications.	64
6.5	Conclusion.	66
7	Experimental Results and Analysis	67
7.1	Data Collection	67
7.2	LCM Results	78
7.3	LCM Accuracy	78
7.4	Model Accuracy	80
7.5	Conclusion.	84
8	The Global Connectivity Map	87
8.1	Relevant Research	89
8.2	Overview	91
8.3	GCM Definition.	92
8.4	Summary	96
9	Building the Global Connectivity Map	99
9.1	Gaussian Process Regression.	99
9.2	Optimal Estimation	105
9.3	Summary	116
10	Refining the Map	119
10.1	Problem Statement.	120
10.2	Recent Research	120
10.3	Overview	121
10.4	Structural Learning in Graphical Models	128
10.5	GCM Inference	129

11 GCM Experimental Results	137
11.1 GCM Results (March 2013)	137
11.2 GCM Results (July 2012)	143
11.3 GCM Measurement Strategies	144
11.4 Conclusion.	145
 12 Conclusions and Recommendations	 155
12.1 LCM	155
12.2 GCM	156
12.3 Model Comparison	158
12.4 Temporal Model Considerations	159
12.5 GCM Techniques	160
12.6 Additional Future Research	161
12.7 Conclusion.	163
 List of References	 165
 Initial Distribution List	 179

THIS PAGE INTENTIONALLY LEFT BLANK

List of Figures

Figure 4.1	<i>A comparison of the covariance and semi-variogram function. The covariance function is drawn in blue and the semi-variogram in red. The plot shows the “inverse” relationship between the two.</i>	25
Figure 4.2	<i>An example of an empirical semi-variogram values (in red) together with a best fitting (using least linear squares) exponential semi-variogram. The plot includes the locations of the semi-variogram nugget, sill and range.</i>	28
Figure 4.3	<i>The h-scattergram of 152 RSS measurements taken from bearing window of 80 degrees. The lag distance (h) is 10 meters.</i>	29
Figure 4.4	<i>Examples of semi-variogram models.</i>	32
Figure 5.1	<i>Simulated acoustic modem RSS measurements (in dB).</i>	36
Figure 5.2	<i>The empirical semi-variogram calculated from the simulated acoustic modem RSS random field.</i>	37
Figure 5.3	<i>An overhead picture of Monterey harbor where boat traffic, piers and moored boats can play havoc on acoustic signals.</i>	38
Figure 5.4	<i>RSS measurements relative to an static buoy collected in Monterey harbor.</i>	39
Figure 5.5	<i>(Left) RSS measurements over a 45-degree bearing window. (Right) The resultant range semi-variogram. The maximum lag-distance for the semi-variogram is 100m. The nugget is approximately 11, the sill 26 and the range 50.</i>	40
Figure 5.6	<i>(Left) RSS measurements over a 45-degree bearing window. (Right) The resultant range semi-variogram. The maximum lag-distance for the semi-variogram is 100m. The nugget is approximately 6, the sill 26 and the range 80.</i>	40

Figure 5.7	<i>(Left) RSS measurements over a 45-degree bearing window. (Right) The resultant range semi-variogram. The maximum lag-distance for the semi-variogram is 100m. Selection of the sill and range is trickier here. The nugget is approximately 4, the sill is 13 and the range is 70, even though the semi-variogram increases after a lag distance of 80. The reason not to select a higher sill/range value is not to have a covariance function that influences the estimate with far away data points.</i>	41
Figure 5.8	<i>(Left) RSS measurements over a 45-degree bearing window. (Right) The resultant range semi-variogram. The maximum lag-distance for the semi-variogram is 100m. The nugget is approximately 1, the sill 5 and the range 50.</i>	42
Figure 5.9	<i>The radial semi-variograms from the S_1 acoustic modem data set collected in Monterey Harbor, July 2012. Each semi-variogram is labeled based on the fixed distance from the source position.</i>	43
Figure 5.10	<i>The radial semi-variograms from the S_2 acoustic modem data set collected in Monterey Harbor, July 2012. Each semi-variogram is labeled based on the fixed distance from the source position.</i>	44
Figure 5.11	<i>The radial semi-variograms from the S_3 acoustic modem data set collected in Monterey Harbor, July 2012. Each semi-variogram is labeled based on the fixed distance from the source position.</i>	44
Figure 5.12	<i>The radial semi-variograms created from the S_4 acoustic modem data set collected in Monterey Harbor, July 2012. Each semi-variogram is labeled based on the fixed distance from the source position.</i>	45
Figure 5.13	<i>The radial semi-variograms created from the S_5 acoustic modem data set taken from Monterey Bay in July 2012. Each semi-variogram is labeled based on the fixed distance from the cell center.</i>	45
Figure 6.1	<i>The plot represents a simulated data set of Received Signal Strength (RSS) acoustic modem measurements. The measurements are taken relative to a centrally located static buoy as a source node. The measurement intensity values range from 1 to 30dB using the MATLAB colormap jet.</i>	52
Figure 6.2	<i>Figure 6.2a shows the collection of points used to create the variogram. Figure 6.2b shows the resulting semi-variogram.</i>	53
Figure 6.3	<i>The figures represent the Kriging mean and variance (on the left and right, respectively). The original input data is shown in Figure 6.1.</i>	56

Figure 6.4	<i>The plots show the results of a Monte Carlo simulation calculating the Kriging error that is the difference between a true model and a model with added random noise. The red line on the right hand side shows the mean value for each bin and the line gives an indication of the rate of convergence.</i>	58
Figure 6.5	<i>Mean and variance estimates of the RSS random field using the additive Kriging methodology.</i>	61
Figure 6.6	<i>LCM Example for Monterey Harbor, CA: The green dot represents a fixed reference point and is the source for all communication x_c. The boundary of the survey area (Ω) is drawn in red. The overlaid orange lattice represents the map resolution. This has been drawn disproportionately large for the graphic. The light blue dot is the centroid of the lattice cell where the mean estimate and variance is calculated but is assumed to hold inside the entire cell.</i>	63
Figure 6.7	<i>An example of decreasing WiFi Received Signal Strength as a distance increases from the Mostofi paper entitled “Compressed Cooperative Sensing and Mapping in Mobile Networks,” [1].</i>	65
Figure 7.1	<i>The annotated positions in the image (S_1, S_2, S_3, S_4, S_5) are the locations of the gateway buoy for the first two rounds of testing (August 2011 and July 2012). For each dataset, the buoy (shown on the right-hand side) was placed at one of the locations and signal statistics were collected using this buoy as the source of a message packet where the destination was from a boat maneuvering around the harbor.</i>	68
Figure 7.2	<i>Description of the differences between the Input SNR, Output SNR and Symbol SNR with the WHOI Micro-modem. Graphic courtesy of Lee Freitag WHOI Principal Scientist.</i>	69
Figure 7.3	<i>The RSS measurements relative to the S_5 buoy position.</i>	70
Figure 7.4	<i>The RSS measurements relative to the S_4 buoy position. On the right-hand side show a magnification around the commercial wharf.</i>	71
Figure 7.5	<i>The RSS measurements relative to the S_3 buoy position.</i>	72
Figure 7.6	<i>The RSS measurements relative to the S_2 buoy position.</i>	73
Figure 7.7	<i>The RSS measurements relative to the S_1 buoy position.</i>	74

Figure 7.8	<i>RSS measurements (relative to S_2 buoy) partitioned into 4 groupings annotated by the red, green blue and magenta triangles</i>	75
Figure 7.9	<i>RSS plots and best fitting polynomial functions corresponding to the red, green, blue and magenta triangles from Figure 7.8</i>	76
Figure 7.10	<i>Tides for July 30, 2012</i>	77
Figure 7.11	<i>RSS measurements relative to the S_1 buoy position. July 2012</i>	78
Figure 7.12	<i>RSS measurements relative to the S_2 buoy position. July 2012</i>	79
Figure 7.13	<i>RSS measurements relative to the S_3 buoy position. July 2012</i>	80
Figure 7.14	<i>RSS measurements relative to the S_4 buoy position. July 2012</i>	81
Figure 7.15	<i>RSS measurements relative to the S_5 buoy position. July 2012</i>	82
Figure 7.16	<i>Kriging mean and variance results from the S_1 dataset</i>	82
Figure 7.17	<i>Kriging mean and variance results from the S_2 dataset</i>	83
Figure 7.18	<i>Kriging mean and variance results from the S_3 dataset</i>	83
Figure 7.19	<i>Kriging mean and variance results from the S_4 dataset</i>	84
Figure 7.20	<i>Kriging mean and variance results from the S_5 dataset</i>	84
Figure 7.21	<i>The image overlays the background noise (dB) on top of a binary representation of the Monterey Harbor. The color scale is the Matlab “jet” colormap.</i>	85
Figure 7.22	<i>The aggregated Kriging error for the datasets (S_1, S_2, S_4, S_5) using the k-fold cross validation where $k=10$. The x-axis is the Kriging errors divided into bins of equal size and the y-axis is the number of elements in each bin.</i>	86
Figure 8.1	<i>The plot gives a visual interpretation of a Riemann sum such that the integral of the function over the range interval $[0,300]$ is approximated by a sum of rectangles such that the upper right point of the rectangle is the RSS approximated by the Matérn function (in red) at the appropriate range. .</i>	94

Figure 8.2	The above graphic displays an overlay that represents a visual depiction of a Global Connectivity Map. The red and yellow borders form the permeable and impermeable boundaries for an undersea GCM. The impermeable boundaries are in red and are examples of where no communication is possible outside the boundary. The permeable boundaries in yellow are examples where communication is possible outside the boundary. The gray grid represents the resolution of the map such that the goal of the estimation process to produce an LCM that is centered on the grid cell (in light blue).	97
Figure 9.1	A selection of Matérn covariance functions.	104
Figure 9.2	The graphic depicts the locations of acoustic signal strength measurements (in blue) from Monterey Harbor and the boundaries of the bearing window (in red) from 165 to 195 degrees where the points within the boundaries are used to estimate the random function.	105
Figure 9.3	The plot displays an estimate of a random function based on 50 S_1 measurements in the bearing window from 165 to 195 degrees.	106
Figure 9.4	The plot displays an estimate of a random function based on 122 S_1 measurements in the bearing window from 165 to 195 degrees.	107
Figure 9.5	The plot displays an estimate of a random function based on 50 S_1 measurements in the bearing window from 165 to 195 degrees using a Matérn prior mean function.	108
Figure 9.6	The plot shows the estimate of a random function based on 122 S_1 measurements in the bearing window from 165 to 195 degrees using a Matérn prior mean function.	109
Figure 9.7	The figure show a dataset of RSS measurements taken in March 2013 in Monterey Harbor, CA. It is partitioned into quadrants north, east, south and west	112
Figure 9.8	The figure shows the north, south, east and west plots of RSS values over increasing range. This is the same dataset the previous figure.	113
Figure 9.9	SCKF Results on the North, East, South and West data sets.	114
Figure 9.10	The convergence of the system covariance. The four plots show the trace of the SCKF covariance matrix for the North (Red), East (Green), South (Blue) and West (Black) data sets.	115

Figure 9.11	<i>On the left is a Matérn function and on the right is a cubic spline. The Matérn function impacts all estimates along the range while the cubic spline has minimal impact to the surround points.</i>	116
Figure 9.12	<i>Gaussian Process Model results on the North (Red), East (Green), South (Blue) and West (Black) datasets.</i>	117
Figure 10.1	<i>On the left is an example three-node MRF where each of the factors have two possibilities between the nodes. On the right is a representative discrete probability space enumerating the possibilities such that each row has an assigned numerical value.</i>	125
Figure 10.2	<i>The above graphic displays a grid overlay that represents a map resolution of a GCM.</i>	130
Figure 10.3	<i>The plots represent simulated RSS acoustic modem measurements. The measurements are taken relative to a centrally located static buoy as a source node. The measurement intensity values range from 1 to 30dB using the MATLAB colormap hot.</i>	134
Figure 10.4	<i>The plot shows the random functions for LCM_1 and LCM_2 given in Figure 10.3. The top row are the LCM_1 random functions for north, east, south, west and the bottom row are the same for LCM_2.</i>	134
Figure 10.5	<i>Accuracy results of the combined SCKF ℓ_1 Regularized Logistic Regression approach with the simulated example.</i>	135
Figure 10.6	<i>Precision results of the combined SCKF ℓ_1 Regularized Logistic Regression approach with the simulated example.</i>	136
Figure 11.1	<i>Tide chart for February 28, 2013</i>	138
Figure 11.2	<i>First day positions for the static gateway acoustic buoy in Monterey harbor.</i>	138
Figure 11.3	<i>Second day positions for the static acoustic gateway buoy in Monterey harbor</i>	139
Figure 11.4	<i>Combined positions of the gateway acoustic buoy where orange annotates locations where data sets were collected on both days.</i>	140
Figure 11.5	<i>Buoy Positions, the number of measurements taken and the start time for the RSS data collection in Monterey Harbor, Monterey CA February 28, 2013</i>	141

Figure 11.6	<i>Buoy Positions, the number of measurements and the start time for the RSS data collection in Monterey Harbor, Monterey CA March 1, 2013 . . .</i>	142
Figure 11.7	<i>RSS Measurements collected relative to the S_{18} gateway buoy position. In black are the designated positions of the acoustic modem gateway buoy.</i>	143
Figure 11.8	<i>RSS measurements collected on February 28, 2013</i>	147
Figure 11.9	<i>RSS measurements collected on March 1, 2013</i>	148
Figure 11.10	<i>A voronoi partition with the static buoy positions acting as the control points of the partition</i>	149
Figure 11.11	<i>A GCM based on March 2013 RSS acoustic modem data collected in Monterey Harbor, Monterey, CA.</i>	150
Figure 11.12	<i>The LCM reference model. The random fields of the GCM are compared to the reference model to determine the map color scaling.</i>	150
Figure 11.13	<i>Results of the K-Fold cross-validation method. The red curve is a best fitting normal distribution with $\mu = 0.1866$ and $\sigma^2 = 3.27$.</i>	151
Figure 11.14	<i>The function in magenta provides the basis for calculating the LCM reference model. It is used to determine the GCM scalar value for each partition.</i>	151
Figure 11.15	<i>RSS Measurements (July 2012)</i>	152
Figure 11.16	<i>The July 2012 GCM.</i>	152
Figure 11.17	<i>Results of the K-Fold cross-validation method for the July 2012 GCM. The red curve is a best fitting normal distribution with $\mu = -0.08$ and $\sigma^2 = 2.40$.</i>	153
Figure 11.18	<i>July 2012 GCM. The colors correspond to general assessments of signal strength and the vectors centered on the central cell node correspond to directional signal strength.</i>	153
Figure 12.1	<i>Source node S_{44}. Data collected on February 28, 2013</i>	159
Figure 12.2	<i>Source node S_{44}. (a) RSS Data collected at 1009 PST on March 01, 2013, (b) RSS Data collected at 1049 PST on March 01, 2013</i>	160

THIS PAGE INTENTIONALLY LEFT BLANK

List of Tables

Table 5.1	<i>The empirical semi-variograms curve fitting results of the four models. The models are evaluated using RMSE. The lowest error is highlighted in bold text.</i>	39
Table 6.1	<i>A table of six points consisting of position and RSS measurement used for the Kriging mean and variance example.</i>	54
Table 6.2	<i>The table shows the results of a Monte Carlo simulation of 1000 runs estimating the error based on increasing number of points. Given an assumed true model and true semi-variogram, it shows the convergence of the Kriging technique as the number of points in a fixed domain increases. . . .</i>	57
Table 7.1	<i>SNR measurements to each destination (August 2011)</i>	69
Table 7.2	<i>Oceanographic data collected from an AUV in Monterey Harbor, CA (August 2011)</i>	74
Table 7.3	<i>RSS Measurements to each destination (July 2012)</i>	75
Table 7.4	<i>The table provides the mean and variance associated with a best fitting Gaussian distribution for the Kriging error (ϵ for the S_1 through S_5 data sets).</i>	85
Table 10.1	<i>First 20 columns of the design matrix</i>	135
Table 10.2	<i>Last 20 columns of the design matrix</i>	135
Table 11.1	<i>GCM data collection strategies</i>	145

THIS PAGE INTENTIONALLY LEFT BLANK

List of Acronyms and Abbreviations

CAVR Center for Autonomous Vehicle Research
LCM Local Connectivity Map
GCM Global Connectivity Map
NCS Network Control System
GPS Global Positioning System
PGM Probabilistic Graphical Model
SLAM Simultaneous Localization and Mapping
BLUE Best Linear Unbiased Estimator
RSS Received Signal Strength
TOA Time of Arrival
AOA Angle of Arrival
GPM Gaussian Process Model
MRF Markov Random Field
CRF Conditional Random Field
NP Nondeterministic Polynomial Time
IQEKF Iteratively-Quantized Extended Kalman Filter
IQMAP Iteratively Quantized Map
AOSN Autonomous Oceanographic Sampling Network
AUV Autonomous Underwater Vehicle
WHOI Woods Hole Oceanographic Institution
POI Point of Interest
RMSE Root Mean Square Error
MSE Mean Square Error
NLS Non-Linear Least Square
GAM Generalized Additive Models
SNR Signal to Noise Ratio
RHIB Rigid Hull Inflatable Boat
PSK Phase Shift Keying
MFD Matched Filter Detector
RFID Radio Frequency Identification
GSM Global System for Mobile Communications

LOOCV Leave One Out Cross Validation
RF Radio Frequency
TDOA Time Difference Of Arrival
SCKF Spatial Communications Kalman Filter
GP Gaussian Process
AI Artificial Intelligence
MLE Maximum Likelihood Estimation
MAP Maximum a Posteriori
L-BFGS Limited-memory Broyden Fletcher Goldfarb Shanno
KNN k-Nearest Neighbors

Executive Summary

The proliferation of wireless devices with cheap, powerful, small computers has created a society that is increasingly interconnected. The accurate estimation and prediction of wireless communication channels holds the promise to improve a wide variety of networking applications used on these devices. This dissertation presents a rigorous mathematical framework that supports efficient construction of wireless connectivity maps with minimal measurements.

In characterizing the signal strength of the communications channel, one can envision strategies that are evaluated against the accuracy, time horizon and complexity of the environment. Current common strategies include a simple signal attenuation model, a physics-based model and a data-driven model. Each has attributes and disadvantages. The signal attenuation model is simple to implement but is inaccurate in more challenging environments. The physics-based model may be more accurate than the attenuation model but requires active sensing and may have limited accuracy in challenging environments.

The third approach is a signal strength survey of the area. At its extreme, it requires a complete channel characterization through extensive measurements. While this is obviously time consuming and inefficient, it does produce an accurate model. With this approach, there is no aspect of predictive mathematical modeling involved since estimates are produced from measurements in the environment.

This dissertation seeks a data-driven solution to build a situational understanding of the communication environment with minimal channel measurements. The goal is to produce a more nuanced understanding of the communication landscape without fully measuring the environment. The fundamental contribution of the dissertation is a methodology for stochastic estimation of communication constraints between multiple agents. This methodology is flexible enough to be used in a variety of environments, but it may be most useful in cluttered, dynamic areas. It is presented in two distinct yet related sections — part one and two. The first part focuses on the Local Connectivity Map and the second part focuses on the Global Connectivity Map.

The Local Connectivity Map (LCM) is a measure of signal strength relative to the local body coordinates of an agent. The LCM is modeled as a Random Field. The technique

to estimate the LCM Random Field is known as Kriging. It is a Best Linear Unbiased Estimator developed by the Geostatistics community. It produces a mean and variance estimate of the Random Field.

The LCM is limited by its ability to provide estimates relative to the local reference frame. We are also interested in providing estimates regardless of the agents position in a bounded region. A new structure called the Global Connectivity Map (GCM) is presented for this purpose. It is a novel methodology designed to provide an estimate of the ability to communicate between any two positions within the survey boundaries. The GCM is typically defined as a two or three dimensional lattice of Random Fields. Each cell within the lattice is a LCM. The size of the lattice corresponds to the GCM map resolution. For each lattice cell, an integral mean and variance estimate is produced for summarizing the ability to regionally communicate. It also produces a signal mean and variance along a discrete number of bearings over increasing distance.

A key design consideration in the development of the framework is the ability to construct the LCM and GCM as the measurements are collected — in near real-time. This is in contrast to the alternative of building the connectivity maps in a batch process after collection is complete. One mathematical impact of this design choice is to combine temporal and spatial optimal estimation techniques in a Bayesian recursive filter.

An application of particular interest is the collaboration of unmanned systems, and in particular unmanned undersea vehicles. Validation of the framework is demonstrated with undersea acoustic communications. This is perhaps the most difficult communication environment for cooperative multi-vehicle operations. This is due to limitations of acoustic signal propagation.

In summary, the fundamental contribution of the dissertation is a mathematical framework for stochastic estimation of communication constraints between multiple agents. In addition, the methodology is tested and analyzed using multiple sets of undersea acoustic modem signal measurements collected over multiple years in Monterey Harbor, Monterey. It is believed that this methodology is flexible enough to be used in a variety of environments and can be used within a networked control system to potentially provide greater flexibility, more robustness and greater energy efficiency.

Acknowledgements

First, I would like to thank the long time support and mentorship of Tom Swean. I have long appreciated his sustained commitment towards science, research and the Navy mission. When he retires, he will be sorely missed. Second, I'd like to thank Tony Healey for his friendship and enduring confidence in me. It has been an honor to work for him and alongside him for the past twelve years and I look forward to the next twenty. Next, I would like to extend my deep appreciation to all members of the dissertation committee — Geoff, Mathias, Joel, Tim, and John. I can not imagine a more talented committee and I am very grateful for your sound advice and persistence through these years. To the members of the NPS Center for Autonomous Vehicle Research, especially Sean, Aurelio, Noel and Tad, it has been my honor to work alongside you; thank you with helping me to achieve this degree. To my wonderful family—my beautiful daughters, Sedona and Delaney, you both have been an inspiration and a joy to watch grow. Julie, I Love You. Hand in hand we storm together through life and I am the luckiest of men.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 1:

Introduction

In the last ten years, we have witnessed a marked improvement in unmanned mobile systems. Autonomous underwater, surface, aerial and ground vehicles are revolutionizing commerce, industry and the military. The vehicles permit a reduction in operating costs and an improvement in productivity for wide varieties of traditionally human-labor intensive activities. Their ability to cover large areas, operate in harsh environments and maintain a persistent presence permits unprecedented monitoring and access. That said, there is a general consensus that significant improvements are still possible.

Three major areas of emphasis are intelligent autonomy, energy efficiency and cooperation. Intelligent autonomy is the ability of the unmanned system to improve its mission effectiveness through a better awareness of the environment. This could be as simple as obstacle detection and avoidance or extremely challenging tasks such as maneuvering through a cluttered and dynamic area where mission objectives change as the environment is modified. Sub-disciplines for developing intelligent autonomous behaviors include four components: perception (sensory input), situational and environment awareness (map building) and motion planning (path planning and path following).

Second, these unmanned systems have limited onboard fuel or battery capacity. This is a major limitation for today's systems. Improvements in how they use limited energy supplies and the development of creative ways for recharging will have important ramifications about their future utility. Improvements in energy efficiency can come about in several ways: 1) Better energy density of battery systems, 2) Better utilization of limited onboard energy, 3) Utilization of latent environmental energy (waves, wind and solar), 4) Autonomous energy replacement systems such as docking stations.

Cooperation amongst the unmanned agents is the third area of emphasis. It has the potential to impact a wide variety of operations through improved search rates, and system robustness. In this case, the robustness of the system is defined, in part, through redundancy (if one agent fails to complete a task, another is in close proximity and can replace it). Cooperative estimation is another aspect for improving system robustness. This is the

process by which each agent can leverage information from others to improve situational awareness. One example is collaborative position estimation and another would be sharing sensory data to improve navigation or the resolution of a collective survey. It is the communication within a cooperative system that is the focus of the dissertation.

1.1 Network Control System

A major trend within industrial, military and commercial systems is the blending of communications within a Networked Control System (NCS) [1], [2]. This takes into consideration of the limitations and strengths of communication links for collaborative control of a multi-agent system. It has been identified as an important developing research topic that impacts disparate fields such as astronomy, traffic management and nanotechnology. It is critically important in the military domain where distributed and collaborative operations can dramatically speed the pace of operations with the intended goal of tactical, operational and strategic advantage.

A NCS is defined by a distributed group of sensors and actuators whereby a closed control loop is used to bound the performance of the system [3], [4]. When the method of communication between the collaborative entities (or agents) is via wireless links, understanding the issues associated with wireless communications, especially with respect to the system control, becomes critical.

The salient issues within wireless NCS include:

1. **Bandwidth-limited channels** — Communication channels have a maximal volume of information that can be distributed per unit time. The Shannon-Hartley theorem describes the maximum rate of information possible over a communications channel with a specified bandwidth in the presence of noise. Of interest, with respect to the control system, is the minimum rate that is required to stabilize a feedback system over a finite capacity channel.
2. **Sampling and delay** — This refers to network latencies that are the result of encoding and decoding communication signals and the propagation delay between nodes.
3. **Packet dropouts** — This refers to the possibility of a packet not being delivered to the intended recipient. This is more common in wireless networks. It could be due to:
 - (a) Unanticipated changes in the network topology due to movement.

- (b) Changes in the communications environment (e.g., movement in a cluttered, dynamic region).
 - (c) Buffer overflows due to network congestion.
4. **Systems architecture** — Especially prevalent within wireless networks is the discovery and maintaining of routes to ensure information is adequately and reliably distributed. Route discovery and maintenance protocols within mobile networks can take significant bandwidth. This becomes critical in networks with limited bandwidth.

1.2 Distributed Coordination of a Wireless NCS

A further sub-discipline of NCS is the distributed coordination of a multi-agent mobile system. This considers control of mobile wireless systems that work in collaboration towards a shared task. A general summary of cooperative control modalities include:

1. **Collaborative surveys** — A team of mobile agents is responsible for area coverage such that each agent is responsible for one or more partitions of the entire assigned survey area. Collectively they cover the complete area. An example of a collaborative survey would be an altitude map of an urban environment with several unmanned aerial vehicles.
2. **Collective sensing** — Mobile agents work collectively to identify a source. This could be an environmental phenomenon of interest or a military target. These are mobile sensor platforms that together, through cooperative control, act as a distributed wide aperture array. The control associated with each agent seeks to maneuver the sensor to maximize the probability of detection with regard to the entire system.
3. **Collaborative work** — Today's state of the art with unmanned system's emphasizes navigation. As we move beyond navigation, future systems will focus on work that can be accomplished by a team of agents. This is exemplified by a team of quadrotors collectively building a structure [5].

The difference in communication between a NCS and a multi-agent mobile systems is the impact of navigation on the wireless communication. The relative motion between vehicles can impact communications in several ways: First, the mobility impacts valid communication paths within the network. This places an emphasis on maintaining routing tables to determine how to route information amongst agents. The routing protocol must

run constantly and can add significant communication overhead.

Second, in some environments there can be uncertainty about the position of the agents. This is especially true in underwater, space and indoor environments where there is no GPS signal. The positional uncertainty can impact the ability to send messages since the difference between the true and estimated position could determine whether communications is possible.

Third, the relative motion between vehicles can impact their ability to communicate. This is especially true for underwater acoustic networks where the ratio between the vehicle velocity and the speed of wireless transmission (underwater speed of sound is nominally 1500 meters per second) is much greater than in a radio wireless network. This means that doppler shifts in the received signal can have a greater impact on correctly receiving messages.

1.3 Feedback Control of Cooperative Systems

Feedback control involves stabilization of system performance through error reduction between a desired system goal and a measurement of its present state. Distributed cooperative feedback seeks system error reduction such that the collective errors associated with the individual system elements are minimized. Designated system goals could include navigational accuracy, maintaining a minimal level of communications or an estimate of an environmental feature.

A natural tool for modeling cooperative control and information flow within the mobile network is algebraic graph theory. Here, the multi-agent system is represented as a graph where the vertices are agents and the valid communication paths between the agents are represented as edges in the graph. A particularly useful framework within graph theory for determining conditions of stable cooperative control is the use of the Laplacian matrix [6].

The Laplacian matrix describes the state of connectedness of the network. It is defined as $L = D - A$ where A is the adjacency matrix for the vertices and the D matrix is the out-degree matrix where entries are limited to the diagonal and represent the number of outgoing network links of each vertex. Much of the research into control of cooperative systems involves the analysis of the controllability of the network using the Laplacian, given various assumptions about the edges in the graph. A common approach within this

literature is to define a Lyapunov stability function which states the conditions under which the system remains stable [7].

One branch of cooperative control is described as a consensus algorithm where the goal is to agree on parameters based on shared information. Information consensus is a means by which to achieve collective group behavior through local interaction. The idea is that for a team of vehicles to work cooperatively towards a shared objective it is necessary that the vehicles have a consistent view of the information that is critical to the task. Examples of a consensus control problem might include the location of a surveillance target or agreed upon locations for communicating between agents, where each agent has a designated role within of a collaborative survey of a region. In sum, this is a useful framework for modeling information flow for different types of coordination tasks such as swarming, flocking, synchronization and distributed surveys.

There are different assumptions about the graph edges that can be used to represent information flow in a mobile network. Graph edges can be directed or undirected to reflect one-way and omni-directional communications. The edges can also be weighted to reflect constraints in the information flow. The consensus problem has been analyzed from the standpoint of a random, weighted, directed graph [8], whereby existence of any edge is assigned a probability. This representation is a reasonable assumption for developing cooperative control using wireless mobile networks.

Ultimately, cooperative control hinges on realistic assumptions about the performance of the network that influences the network connectedness. It would be valuable to replace the assumptions with actual stochastic estimates about the random graph. This is the goal of the dissertation—to provide better control of cooperative multi-agent deployments through better understanding of the communications channel. This better understanding of the communications environment can be used to determine probabilistic values assigned to the random graph to provide better performance for consensus-based (and other) cooperative algorithms.

1.4 Communication Channel Estimation

There are different approaches for communication channel estimation. One can envision the choices as a continuum along a triangular plane. On one side is a pure physics-based modeling approach. The goal is to capture an accurate ability to communicate through

physical modeling. It involves characterization of the communications channel through a predictive mathematical function. It accounts for the vagaries of the communication medium (air, water, space) and the nature of the signal (acoustic, radio, light) to produce an estimate of the signal between one agent attempting to communicate with another. It requires no measurements of the channel.

On another side of the continuum is a signal attenuation model. It is a representative function for estimating signal strength with respect to distance. The strength of the approach is its simplicity, the downside is that it may be particularly inaccurate in difficult communication environments. On the third side of the continuum is where the estimation process is fully determined through sensor measurements. The strength of this approach is that it requires no physical modeling. This has particular utility in cluttered, dynamic areas where propagation of signals can result in unusual but effective communication paths. The weakness in the approach is that it is difficult to fully measure the environment with limited amount of time, energy and resources.

This dissertation seeks a solution in-between the extremes of the continuum—to build a situational understanding of the communication environment with less channel measurements. The goal is to produce a more nuanced understanding of the communication landscape without fully measuring the environment. It is hypothesized that in combination with a feedback control system it would produce a more robust and effective system with increased navigation flexibility and improved energy efficiency.

1.5 Main Hypothesis

An improved understanding of a communication channel is possible through a data-driven approach with a limited number of sensor measurements. This characterization of the communication channel can be potentially used in a variety of domains including collaborative navigation amongst a system of unmanned vehicles. Because of the difficulties associated with communications, especially in cluttered and dynamic environments, this is a potentially better approach than alternatives such as physics-based and signal attenuation modeling. Furthermore, it is possible to estimate the communications channel in near real-time operations (as opposed to post-processing collected data) and gain an understanding with respect to the signal uniformity over a bounded region.

1.6 Contribution of Dissertation

The contributions of the dissertation are four-fold: First, a hierarchical approach for wireless communication signal strength estimation. Second, is the development of a recursive filter for real-time communication channel estimation. Third, is the use of structural learning in graphical models for evaluating signal similarities between regions. Fourth, is demonstration of the techniques through the collection and analysis of a significant under-sea acoustic communication dataset. The common mathematical framework for describing these techniques is the theory of random fields for spatial estimation.

Regarding the first contribution, the hierarchical approach consists of two tiers, the Local Connectivity Cap (LCM) and the Global Connectivity Map (GCM). The LCM seeks to build a Random Field based on limited signal measurements relative to a local body reference frame or equivalently a single, fixed reference point. The GCM is a lattice of LCMs that can be used for estimating signal strengths between any two agents in a defined survey region.

The second and third contributions are techniques that are employed to enable near real-time estimation of the GCM. The Bayesian recursive spatial filter permits sequential spatial estimates of the Random Field (or function) as measurements become available. One can potentially reduce the number of measurements if regions display similar signal characteristics. A technique is presented for comparing regions of communication using Probabilistic Graphical Models (PGM) and structural learning of network graphs.

The third contribution is a substantial collection of acoustic modem signal statistic measurements in Monterey Harbor, Monterey, CA. The undersea domain is a difficult environment for communications. This is even more so in a cluttered, dynamic undersea region such as a harbor. The dataset highlights the sometimes unusual signal propagation paths prevalent in these areas.

1.7 Organization of Dissertation

The dissertation is structured in two distinct, yet related parts. Part one emphasizes the Local Connectivity Map (LCM). It contains four chapters. The first describes structural analysis of spatially-related variables. This is the process of building a semi-variogram which is key for stochastic spatial estimation. The second chapter describes aspects of structural

analysis that are unique to a radial broadcasting source like a communication transmitter. The next chapter discusses Kriging, again with an emphasis on communications. The final chapter discusses the application of the theory to a collected dataset of acoustic modem measurements taken in Monterey Harbor, Monterey, CA. This chapter includes analysis of the developed techniques.

Part two emphasizes the Global Connectivity Map (GCM). It contains five chapters. It starts with a chapter on defining the GCM. The second chapter discusses the theory associated with building the GCM. The third chapter continues with techniques for simplifying the GCM by introducing structural learning for networks in probabilistic graphical models. The fourth chapter applies the techniques developed in the previous two chapters toward another acoustic modem dataset again collected in Monterey Harbor. The fifth chapter provides analysis of the experimental results. Before getting started with part one, there is a chapter that provides a review of past literature. After part two, the dissertation concludes with final remarks and recommendations for future work.

CHAPTER 2:

Background

This chapter reviews past research related to the topic area. The approach taken is necessarily multi-disciplinary in nature and leverages techniques developed from the fields of statistics, algebraic graph theory, networking, robotics and machine learning. A review of collaborative navigation and undersea acoustic networking literature is also included. The research review is broken down into two sections that relate to the LCM and GCM. Each reflects a robotic-centric communications approach to improve multi-vehicle collaboration and coordination.

2.1 Local Connectivity Map

The problem that is addressed in the first part of the dissertation is the estimation of signal strength relative to a local reference frame. For static, immobile nodes one might be interested in determining where they might be placed to maximize an objective function such as coverage or throughput. For mobile nodes, estimation of signal strength relative to the pose of the robot would be useful for determining constraints in formation control.

There are several techniques available for estimating signal statistics. These include regression methods, Artificial Neural Networks, Radial Basis Functions, splines and statistical interpolation. Of particular interest is the latter. Principally, the reason is the value of the variance estimate in a stochastic framework. In an information theoretic context, it helps determine how a cooperative navigation algorithm might optimize path planning considerations given high (or low) levels of communication uncertainty.

Given a preference for a stochastic approach, the building of the Local Connectivity Map leverages the theory of random fields for signal (or channel) estimation. A random field is a set of random variables that are correlated in space. In other words, knowledge about a random variable at a particular location gives one the ability to infer about the value of another variable located nearby. This is the essence of spatial statistics. Random fields are of particular interest to the environmental sciences where the goal is to provide density estimates of a regional variable of interest. Standard monographs for the theory of random fields include Vanmarcke [9] and Adler [10] for geometric random fields and standard texts

in spatial statistics include [11], [12], [13] and [14].

Since it is difficult to completely measure the environment for a particular variable of interest, interpolation techniques are frequently required to infer about mean estimates and variance at each position within the random field. The techniques that are used within the geostatistics community are sometimes collectively called Kriging [15]. It is named for a South African mining engineer George Krige who was influential in the initial development of the techniques [16]. An important step in Kriging is determining the spatial correlation between measurements. This is known as variogram analysis [17]. Besides use in Kriging techniques, it has recently been used for image classification for remote sensing applications [18].

Part of variogram analysis is determining a proper model that fits the data. Choices include parametric and nonparametric models. There are many examples of parametric models used for semi-variogram analysis. These include spherical, Gaussian, Sine Hole, exponential and Cauchy model, to name just a few. Stein [19] and Haskard [20] discuss the attributes of a parametric Matérn spatial covariance model especially as it relates to anisotropic data. The ability to estimate an anisotropic random field is useful since it provides potentially greater insight into the communications channel as a function of range and bearing. This is especially useful within cluttered and dynamic environments.

Non-parametric models such as splines and calculus of variations methodologies have utility in the function's flexibility or elasticity [21]. The use of a non-parametric model for semi-variogram (and covariance) analysis to capture nonlinearities (that could otherwise not be adequately "handled" by parametric models) makes the implicit assumption that the variability in the data is an accurate representation of the true spatial relationship.

Recent research includes non-parametric estimation using a nearest neighbor estimator with a non-constant smoothing field [22] and the development of a nonparametric technique using a spectral representation [23]. Spectral analysis requires a gridded survey [14]. In this particular application, this is not a preferred methodology since the collection of signal strength measurements is not the primary mission focus (the primary focus is a collaborative survey). Maintaining flexibility in the vehicle's navigation to completely cover an area is an important consideration.

With regard to the use of Kriging for communication estimation, Umer et al. presents a distributed Kriging technique for estimating coverage holes in stationary wireless sensor networks [24], [25]. Additionally, there are other interpolation techniques available, these include Inverse Distance Weighting [26] and Radial Basis Functions [27]. The main advantage for Kriging is that through the semi-variogram modeling it provides measures of map uncertainty (variances).

Wireless localization is a related field where the goal is to determine the location of a mobile device from its wireless signals. A discussion of ultrawide bandwidth signals and the fundamental limits of ranging accuracy is given in [28], [29] and [30]. A survey of methodologies for indoor wireless localization is given by Gholami [31]. Larkin introduced the idea of a communications map for wireless signal strength in mobile adhoc networks. The approach focuses on identification of areas with a particular signal loss [32].

In [33] and [34], Cottingham describes algorithms for processing large amounts of signal strength data for vehicle wireless coverage maps. Pogel and Wolf develop connectivity maps based on analysis of 3G network characteristics for vehicular traffic [35]. Additional research into radio frequency localization techniques include [36], [37], [38]. Recent surveys of underwater localization techniques are given by [39], [40].

Simultaneous Localization and Mapping (SLAM) is a technique for position localization through the identification of sensory detected features [41], [42]. Normally SLAM is used in conjunction with cameras and sonars. These techniques can also be applied within the context of radio signals. Ferris et al. [43], use a Gaussian Process Latent Variable Model in combination with a novel WiFi-SLAM technique [44] for building wireless signal strength maps. These are maps relative to a stationary WiFi access point. Gutmann et al. present Vector Field SLAM—a localization technique focusing on estimating the parameters of a piecewise linear function [45].

It is worthwhile to recognize the similarities between the Kriging techniques used in the geostatistical community and the Gaussian Process Models (GPMs) popularized more recently by the artificial intelligence and robotics communities (among others). At its core, GPMs use the same mathematics as the Best Linear Unbiased Estimators (BLUE) associated with Kriging techniques. The differences are based on the formulation or representation of the problem and the pertinent application.

An area of recent research activity has been the use of compressive sensing techniques together with wireless networks. In compressive sensing, if a signal can be represented in a particular basis function (such as a fourier or wavelet basis function) it has been shown that the signal can be almost perfectly recovered with significantly less sampling than conventional Nyquist/Shannon limits dictate [46]. Mostofi has leveraged the technique for compressive mapping in the estimation of communication signal strength for a sparsity-based communication channel prediction [47] and for cooperative spatial mapping [48]. Additional work attempts to use active learning techniques to further reduce the sensing requires within the compressive sensing framework [49], [50] and [51].

A survey of cooperative localization is given by Patwari (et al.) [52] It reviews some of the limitations to localization that are also pertinent to communication channel estimation. These parameters include Received Signal Strength (RSS), Time of Arrival (TOA) and Angle of Arrival (AOA). Research by Strom and Olson describe signal strength prediction for a team of ground robots which includes improving estimates by leveraging the robots' LIDAR data [53]. Fink and Kumar present a Gaussian Process Model (GPM) to build an on-line model of signal strength for an source location estimate [54]. Zickler and Veloso present a probabilistic approach for multi-vehicle localization via RSS signals [55].

2.2 Global Connectivity Map

The use of Kriging is based on the assumption that the underlying random field is Gaussian and the covariance function is exactly known through data analysis. When dealing with the estimation of a single random field this is advantageous since one can be deliberate with regard to the selection of data for variogram analysis.

With the Global Connectivity Map we are interested in estimating multiple random fields simultaneously. To make this more efficient and “autonomous,” we assume a particular variogram function and estimate the parameters of the function. This is known as Bayesian Kriging. It is useful since it explicitly accounts for the uncertainty associated with the model.

The first contributions for Bayesian Kriging can traced back to Omre [56], Kitanidis [57] and Omre and Halvorsen [58]. Omre attributes the development of Bayesian Kriging as an extension of linear Bayesian theory to spatial problems [59]. Note that in 1984 Kulka-rmi published a technique known as Bayesian Kriging but his techniques have a different

methodology and have subsequently been called “soft-Kriging” [60].

Pilz and Spock [61] provide a useful more recent account of the history and implementation of Bayesian Kriging methods. It highlights and contributes to an important consideration with any Bayesian implementation—how to define the prior distribution. A common assumption is a non-informative prior [62]. Works by Handcock and Stein [63] and Gaudard [64] evaluated the sensitivity of the predictive distribution on prior models. Investigations into new methodologies for defining the prior distribution include the use of a reference prior [65], [66]. Paulo compares non-informative priors with reference priors and Jeffrey’s rule priors [67].

A Bayesian Kriging approach requires the selection of a flexible class of correlation functions. Stein [15] and Haskard [20] both advocate the use of the Matérn family of correlation functions which is based on a modified Bessel function. In particular, Haskard describes the Matérn functions as being particularly useful for supporting anisotropic covariance functions. It can do this through the application of a transformation matrix to the Matérn function.

It is useful to recognize the similarity between Kriging and Gaussian Process Models. GPMs have been influential to the Artificial Intelligence Machine Learning community to address supervised learning problems such as classification. The task typically has input measurements and the goal is to map these to an output vector which classifies the input in an particular category. Examples include image segmentation and natural language processing [68]. Recent work in Bayesian Kriging has focused on the spatio-temporal and large dataset modeling. Katzfuss reflects both of these trends with the publications [69] and [70]. In the latter, the application is the estimation of global CO2 levels.

2.3 Structure Learning with Markov and Conditional Random Fields

An assumption with GCMs is that they can be estimated more quickly through a pre-survey of the mission area. The goal of the pre-survey is to roughly characterize the local random fields and look for regions of similarity with the assumption that if two regions have similar Random Field characteristics future SNR can be used to estimate either field. This greatly reduces the required measurements.

The problem can be posed as a graph structure learning problem where each random field is a graph vertex and the edges represent the conditional dependence between the vertices. In other words, if there exists no edges between two vertices they are conditional independent—knowledge about one vertex provides no information about the second vertex. By collecting signal measurements during the pre-survey one can learn connectivity in the graph and this can be used during the survey so that a single signal measurement can be used for multiple estimations of random fields.

The basis for Structure Learning is PGM. It provides an appropriate framework for analysis of potentially dependent random variables. An excellent reference for PGMs is by Koller and Friedman [71]. Factor graphs [72] represent an important “master class” of graphical models and were introduced by Kschischang et al. [73] and extended by Frey [74].

This extension details the close association of factor graphs with Markov random fields [75] and Bayesian networks [76]. For the problem at hand, we are specifically interested in Markov Random Fields (MRFs). It is a PGM which uses undirected edges in the graph. The use of MRFs makes sense in the context of the problem of interest—it correctly reflects the fact that the conditional dependency works in both directions of an edge between two vertices.

Dogandžić et al. presented a hidden MRF framework for distributed signal processing for sensor networks [77]. It includes a calibration method for estimating model parameters from training data. Also of interest is a subclass of MRFs known as Conditional Random Fields (CRF) [78]. Introductory references for CRFs include [79] and [80]. CRFs model the conditional distribution $P(Y|X)$ where \mathbf{X} is normally a discrete vector of inputs (measurements) and \mathbf{Y} represents a discrete number of classification objects. CRFs offer computational and statistical advantages over generative models that model the joint distribution ($P(X, Y)$).

Early graph structure learning research include the classic algorithm by Chow and Liu [81]. In 2003, Srebro proved that, in general, structure learning for MRF’s is NP-hard. One prevalent method for structure learning is known as ℓ_1 -Regularization. The technique encourages a sparse graph structure by inducing an ℓ_1 -norm penalty on the parameters of the model [82], [83] and [84]. An alternative approach is the use of thin junction trees with low tree width for exact inference [85].

2.4 Cooperative Navigation

The aim of the dissertation is to provide a methodology to estimate the communications channel. The approach is especially relevant to multi-agent mobile systems. A team of cooperating agents can potentially use this information for helping to determine several important questions that pertain to cooperative navigation:

- How to minimize system navigational errors.
- Where to go to improve an agent's probability of successful communications with another agent.
- Where to go to improve the connectivity map.
- How to partition a survey region so that agents conducting searches within these sub-partitions have a near optimal ability to communicate with other agents conducting the survey.

Towards that goal, what follows is a brief review of techniques in cooperative navigation emphasizing the role of communications in the control of multi-agent systems. They are organized in terms of distinct recent contributions towards channel estimation, stochastic channel modeling, location-aware wireless localization, information or network flow and undersea cooperative navigation.

Monographs on cooperative and distributed control of multi-agent systems include the following [86], [87] and [88]. A good overview of communication channel estimation issues for mobile wireless ground systems is provided by [89]. Furthermore in [47], she reviews techniques in the modeling and characterization of wireless channels, this includes the use of compressive sampling theory for signal estimation.

With stochastic channel modeling, the dynamic links in the network topology are assigned probabilities of successful transmission based on position, motion and environmental considerations. This dissertation fundamentally selects a probabilistic view with regard to the modeling of the communications channel. We believe this is warranted given the uncertainty of the signal and the scale of the environments. Porfiri and Stilwell [8] viewed the consensus problem (agreement between a team of mobile agents) within the context of a stochastic information network. Fink et al. [90] adopts a stochastic model that has distinct cyber and physical components to achieve robust control and communication guarantees.

A sub-discipline within cooperative navigation research is distributed multi-vehicle localization. In general, the goal is to share positional information amongst the fleet in order to jointly estimate all robots' poses and minimize the total positional error. This is especially important in GPS denied environments such as space, underwater and indoors. Wymeersch et al. provide an overview on cooperative localization for wireless networks especially with regard to location awareness [91]. Trawny et al. discuss cooperative multi-robot localization under communication constraints and present solutions in terms of a Iteratively-Quantized Extended Kalman Filter (IQEKF) and Iteratively Quantized MAP (IQMAP) estimator [92]. Zhou and Roumeliotis address the problem of determining a two-dimensional pose for pairs of networked robots [93].

Part of the networking problem with multiple mobile agents is determining valid routing paths. Constant messaging passing for a routing protocol can occupy significant bandwidth. For this reason it is a good idea for mobile agents to consider cooperative navigation strategies for minimizing routing protocols and maximizing information flow. Zavlanos et al. [94] considers mobility and routing for control of networked robots. It defines network integrity as the routes and communication rates that are optimal operating points of the network. Ny et al. [95] presents a prime-dual optimization algorithm simultaneously satisfying the navigation task while ensuring desired network communication flow with a base station. Furthermore they generalize this approach for adaptive deployments of mobile vehicles using a stochastic gradient approach [96].

With regarding to undersea navigation with multiple mobile vehicles, Autonomous Oceanographic Sampling Network (AOSN) provided an unique opportunity to test cooperative control concepts for adaptive sampling for a group of autonomous underwater gliders in Monterey Bay, CA [97] , [98] and [99]. This work was unique, in part, from the AOSN approach to a large scale, system oceanography control feedback loop whereby environment models help to determine recommended regions for Autonomous Underwater Vehicle (AUV) deployments to iteratively improve environmental understanding. Kantarci et al. [100] and Han [40] provide survey papers for localization techniques used with underwater acoustic sensor networks. Bahr et al. [101] describes cooperative localization for underwater vehicles through the use of a moving baseline. He also describes the important consideration of inconsistent or overconfident estimates that are possible with multi-agent collaboration when measurements are used multiple times by an agent.

Finally publications by Francesco Bullo, Jorge Cortés and collaborators were influential in this dissertation. These works focus on geometric considerations for coordinated distributed control and are particularly appropriate for distributed survey problems. Specifically, Cortés paper on “Distributed Kriged Kalman Filter for Spatial Estimation” [102] was helpful in recognizing the utility of combining geostatistical tools with control techniques for multi-agent systems to provide potential better insight into spatial processes through random field estimation.

2.5 Undersea Acoustic Networking

The operational exemplar used throughout this dissertation involves underwater acoustic networking. More specifically, the motivating application is a collaborative survey in a confined harbor where the goal is to keep the vehicles underway for long periods without surfacing to prevent navigational hazards. Navigation without GPS requires the aforementioned SLAM approach.

Acoustic communications in the ocean presents its own unique challenges for mobile networking [103] and [104]. This is intrinsically linked to the nature of signal (sound) and the medium (the ocean). Surveys of relevant issues in underwater networks include [105], [106] and [107].

Recent work involving the statistical characterization and modeling of acoustic communication channels include the following references [108], [109], Stefanov and Stojanovic [110] analyze the simulated performance of ad-hoc networks with interference. The Woods Hole Oceanographic Institution (WHOI) micro-modem was used for acoustic testing a data collection. Modem description and a discussion of performance is available here [111], [112] and [113].

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 3:

The Local Connectivity Map

The central focus in part one is the ability to build a Local Connectivity Map (LCM). This is exemplified in our everyday lives when we walk into a coffee shop with the intent of surfing the internet. With prior accurate knowledge of the map one can make a deliberate choice as to where to sit. Better choices result in higher signal strength and, through dynamic rate scaling (part of the IEEE 802.11 wireless standard), it results in a better throughput rate between the coffee shop's access point and the user's laptop.

Wireless communications can also take curious paths between a source and destination. It can reflect, penetrate and be absorbed by various surfaces in unpredictable ways. A mathematical model that predicts signal strength based solely on physical modeling of signal and environmental characteristics is frequently insufficiently accurate to predict channel strength. This is particularly true in cluttered and dynamic areas.

Standing at a distance and bearing from an access point, one can measure the signal strength. Over time, remaining at the same position, the measurements can be summarized with statistics such as the mean and variance over the samples. If one continued to conduct signal measurements by moving and collecting data at surrounding locations, the result would be a set of statistics with regards to a field of random variables. This is known as a random field. Dealing with the relationship between these random variables as it pertains to position is known as spatial statistics.

To begin, we are interested in building a signal model relative to a local reference frame. That is, an estimate of the ability to communicate relative to the pose (location and orientation) of a robot. The robots could be stationary or in a mobile formation but the important point is the estimate is relative to the robot's pose—this is what is meant by a local reference frame.

If each robot had an accurate local connectivity map, one could assign positions to the robots based on an optimization function for maximizing some desirable objective such as coverage or wireless data throughput. For mobile robots, this information could be used in the same way for improving formation control where the objective, for each robot, is to

maintain a pose relative its neighbors.

The LCM is not a panacea for all types of collaborative control. It is not adequate for signal estimation between arbitrary points because the model is relative to the pose of the robot. What is required is a construct that supports signal prediction regardless of positions within a bounded region. This could be used for a variety of applications.

For example, in areas with poor signal strength, robots would use the signal model to remain closer together in order to maintain full communication connectivity and conversely in areas with greater signal strength they would have more latitude/freedom to conduct tasks while maintaining full connectivity. The ability to assess an environment for communication efficacy is the focus of the second half of the dissertation and is called the Global Connectivity Map (GCM). It is global, since the map is defined with respect to a global reference frame.

The structure of part one includes the following: Chapter Four deals with structural analysis. Of principal interest, is a data-driven model for understanding the spatial relationship between variables. In the Geostatistics community, this is known as variogram analysis. It is used to construct a semi-variogram and provides a critical input to the interpolation technique.

The LCM is unique in that it is possible to conduct polar or spherical variogram analysis by considering spatial relationships in both range and bearing. Chapter Five describes the results of building a covariance function in bearing and the impact that it has on the estimation of the LCM. Chapter Six presents interpolation techniques collectively known as Kriging. Chapter Seven presents LCM results and analysis using underwater acoustic data collected in Monterey Harbor, Monterey, CA.

CHAPTER 4:

Structural Analysis

Of interest is the ability to stochastically estimate the signal strength from a limited number of measurements. There is a point of interest (POI) and we are interested in determining an estimate at that point. There are measurements in the neighborhood that are available to help. The simplest approach would be to calculate the estimate as the mean of the neighboring measurements; but this does not take into consideration the distance of each point from the POI. A weighed average, where the weightings are calculated based on the distance from the POI, could be used to calculate the estimate. But this does not take into consideration the variability of the measurements that were used.

The eventual methodology will use the distance and variability of the data in the calculation of the estimate. But, to start, we need to understand the variability of the measurements. This chapter focuses on how to use a collection of measurements to determine the spatial relationship of data.

A Gaussian random process is defined by its two moments, a mean and variance. A random field is an extension of a random process applied to a domain D that is a fixed subset of \mathbb{R}^d . For most practical applications, the domain will be limited to $d = \{2, 3\}$. It is defined as follows:

Definition 1.1 *Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a complete probability space where Ω is a sample space, which is the set of all outcomes, and ω ($\omega \in \Omega$) corresponds to a particular sampling. \mathcal{F} is a set of events and each event is a set with zero or more outcomes and \mathbb{P} is the assignments of probabilities to events. Then a random field is defined as $\{Z(x, \omega) | x \in D; \omega \in \Omega\}$.*

The random field is frequently simplified to $Z(x)$, where x represents a point. For example such when $d = 2$, $x = \{[x_{north}, x_{east}]\}$ or equivalently a representation in polar coordinates $x = \{[x_{range}, x_{bearing}]\}$ when the point is relative to a known reference point. The value of a random field at location x is $Z(x)$ and $z = z(x)$ denotes an observed measurement or realized value of the random field. The term random function is solely used for $d = 1$, for example the estimate of Received Signal Strength (RSS) as a function of range.

Many of the structural techniques that will be described have been developed by the Geostatistics community where estimation of mineral deposits and oil fields (and many other areas) can have great economic and political impact. G. Matheron defined Geostatistics as “...the application of the formalism of random functions to the reconnaissance and estimation of natural phenomenon” [114].

It is somewhat unconventional to apply Geostatistical techniques to communication channel estimation. But the electromagnetic energy of wireless radio signals are part of our physical natural universe and seen in this light, the use of Geostatistical techniques for estimating the channel is consistent with the original intent of the field.

To determine the intensity of a field it is necessary to take measurements. In the mining industry, the measurements could take the form of core samples; drilling into the land and retrieving a vertical sample of soil. To cartographers it could be a series of altitude measurements of a mountain range. For climatologists it might be the level of CO_2 in the atmosphere.

In all these cases and many others, it is frequently not possible to fully survey an area; only a limited number of core samples can be taken. From those samples, it is necessary to reconstruct an estimate of the variable of interest. Two important tools used in Geostatistics for spatial estimation are variogram analysis and kriging. Together they form the basis for an optimal spatial interpolation technique. These are the techniques that will also be used for building the Local Connectivity Map (LCM).

4.1 Stationarity

Assume for a moment that it is only possible to make a single measurement of a random variable $Z(x_i)$ at a particular location x_i . This provides little information about the probability distribution of this random variable. Suppose it was possible to take additional measurements at other locations $(x_j \mid x_j \neq x_i)$. In order to do statistical inference on a spatial random variable $Z(x_i)$, it is assumed that both x_i and x_j are realizations of the same stochastic process. This is known as the hypothesis of stationarity. With this assumption, two measurements $z(x_0)$ and $z(x_0 + h)$, where h is a positive distance, are now assumed to be two realizations of the same random variable $Z(x_0)$.

There are several kinds of stationarity, three of which are relevant for spatial statistics:

- **Strict stationarity** — A random field is strictly stationary when its spatial law is invariant to translation (represented by h).

$$p\{Z(x_1) < z_1, \dots, Z(x_k) < z_k\} = p\{Z(x_1 + h) < z_1, \dots, Z(x_k + h) < z_k\} \quad (4.1)$$

- **Second-Order stationarity** — A random field is second-order stationary when
 - The expectation of the random field is a finite constant.

$$E[Z(x)] = m, \forall x \quad (4.2)$$

- The covariance $C(h)$ between all pairs of random variables exists and depends not on the location (x), but only on the distance separation (h).

$$C(h) = \text{Cov}(Z(x), Z(x+h)) = E[Z(x), Z(x+h)] - m^2, \forall x \quad (4.3)$$

- **Intrinsic hypothesis** — A random field $Z(x)$ is intrinsic when:
 - The expectation of the random field is a finite constant.¹

$$E[Z(x)] = m, \forall x \quad (4.4)$$

- There exists a finite variance which does not depend on x

$$\text{Var}[Z(x+h) - Z(x)] = E[Z(x+h) - Z(x)]^2 = 2\lambda(h), \forall x \quad (4.5)$$

Note that second-order stationarity implies intrinsic stationarity but not the converse.

For many applications (including using random fields for communication channel estimation), assumptions regarding second order and intrinsic stationarity are relaxed. Specifically, the assumption of a constant mean is unrealistic. Instead, there is an assumption of a constant mean over a bounded region or neighborhood (s) whereby the stationarity assumption holds for locations inside the neighborhood $|h| \leq s$. This is known as *quasi-stationarity*. We will assume the hypothesis of quasi-stationarity for this dissertation unless otherwise indicated. This can be thought of as a compromise between the region of

¹Also frequently stated as $E[Z(x+h) - Z(x)] = 0, \forall x, h$.

homogeneity and the inability to fully measure an environment (especially at a single location) [12].

4.2 Structural Analysis

Structural or variogram analysis determines the spatial relationship between data at different distances. It was originally introduced by Kolmogorov in the 1940's in the study of turbulent flow [115], [116]. There are two steps: The building of the empirical variogram and selection of a function and parameters to accurately represent the data. We start with the definition of the variogram.

4.2.1 Variogram Definition

The variogram "... is defined as the variance of the difference between field values at two locations ($Z(x_1) - Z(x_2)$) across the realization of the field" [11]. Mathematically, this is defined as:

$$2\gamma(x_1, x_2) = \text{Var}[Z(x_1) - Z(x_2)] \quad (4.6)$$

For the more general case, replace (x_1, x_2) with $(x, x + h)$

$$2\gamma(x, x + h) = \text{Var}[Z(x + h) - Z(x)] \quad (4.7)$$

The variance is defined by the expected value of the difference of mean values squared and results in:

$$\gamma(x, x + h) = \frac{1}{2} E[Z(x + h) - Z(x)]^2 \quad (4.8)$$

When we assume second-order (therefore intrinsic) stationarity, the following two relationships hold:

$$\text{Var}[Z(x)] = E[(Z(x) - m)^2] = C(0) \quad (4.9)$$

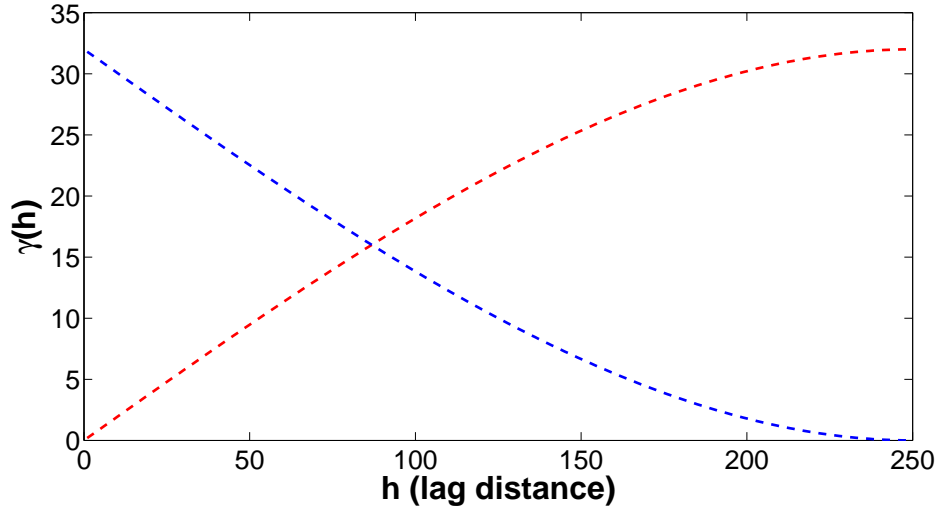


Figure 4.1: A comparison of the covariance and semi-variogram function. The covariance function is drawn in blue and the semi-variogram in red. The plot shows the “inverse” relationship between the two.

$$\gamma(h) = \frac{1}{2}E[(Z(x+h) - Z(x))^2] = C(0) - C(h) \quad (4.10)$$

where C is a covariance function and $C(0)$ is the peak covariance when the distance between variables is zero. In other words, the variance of the random variable is the covariance of two random variables when the distance (h) between the two random variables is zero.

This mathematically demonstrates the “inverse” relationship between the variogram and the covariance. Graphically this is shown in Figure 4.1. The blue dotted line shows the covariance function such that as the lag distance increases the covariance decreases and there exists a distance s where the covariance function evaluates to zero for any distance greater than s . Conversely, the red dotted line shows a variogram function such that beyond the lag distance s it reaches a constant value (the sill). The variogram is described as quantifying the differences between variables separated by a lag distance (h) while the covariance is a measurement of similarity between variables. The relationship between the variogram and covariance is:

$$2\gamma(x_1, x_2) = C(x_1, x_1) + C(x_2, x_2) - 2C(x_1, x_2) \quad (4.11)$$

With the simple relationship that has been shown between the variogram and the covariance one may rightfully ask—What is the utility of the variogram? The answer is two-fold: First, covariance requires the calculation of the data’s mean. This can introduce a bias that doesn’t occur with variograms. Second, covariance is defined with respect to stationary random fields and variograms are defined with respect to intrinsic random fields. This implies that variograms are a more general tool since stationary random fields are a subset of intrinsic random fields.²

An important property for covariance functions is that they must be positive definite. Consider any linear combination of the random function $Z(x_i)$ over N terms $\sum_{i=1}^N \lambda_i Z(x_i)$. Its variance must be greater than or equal to zero.

$$\text{Var} \left[\sum_{i=1}^N \lambda_i Z(x_i) \right] = \sum_{i=1}^N \sum_{j=1}^N \lambda_i \lambda_j C(x_j - x_i) \geq 0 \quad (4.12)$$

The complement definition of valid variogram functions state that it must be conditionally negative definite:

$$\text{Var} \left[\sum_{i=1}^N \lambda_i Z(x_i) \right] = \sum_{i=1}^N \sum_{j=1}^N \lambda_i \lambda_j \gamma(x_j - x_i) \leq 0 \quad (4.13)$$

Armstrong and Jabin [117] show that negative variance is possible when the variogram is not valid.

4.2.2 Isotropy and Anisotropy

A second order random field is isotropic if the spatial relationship between any two points is dependent only on the distance between the points. It is anisotropic if the spatial relationship depends not only on distance but on the angular orientation between the two points. In cluttered and dynamic environments an anisotropic assumption can be useful since changes in spatial relationships as a function of bearing may provide indications of objects or phenomena that are inhibiting communication.

²See Chilés and Delfiner Geostatistics: *Modeling Spatial Uncertainty* 2012, p. 32.

The modeling of an anisotropic covariance function is difficult. A common technique is the development of a transformation matrix such that when it is applied to an isotropic model it converts to an anisotropic model [20]. This approach has utility in many circumstances but may be limited in its ability to describe more challenging anisotropic behavior. For example, with underwater acoustic communications it is not unusual to have small pockets or regions with significantly better signal strength that is a considerable distance from the source and a transformation from an isotropic model would be difficult.

Another approach for modeling anisotropic behavior is to make an isotropic assumption over a limited bearing window such that the sum of all the windows is 360 degrees. This partitions the data such that a variogram can be built that is unique to the bearing window. The anisotropic behavior of the random field is captured through the collective models. This technique is useful for modeling signal statistics, since in difficult environments, signals will not degrade uniformly over distance, but will vary according to bearing. The down side is that it is difficult to determine where to optimally partition the data into bearing windows.

4.2.3 Empirical Variogram

Start with the assumption of second-order stationarity. Consider two measurements $z(x)$ and $z(x+h)$ such that x is a position vector and the measurements are separated by a distance h . The variability between the two measurements is known as the variogram $2\gamma(x, h)$ (and the semi-variogram is $\gamma(x, h)$) and is defined as the expectation of the difference squared

$$2\gamma(x, h) = E[Z(x) - Z(x+h)]^2 \quad (4.14)$$

In practice, the data is divided into bins where the distance between the points are within a bounded neighborhood and within this region the quasi-stationarity hypothesis is assumed to hold. The binned data is evaluated using the following formula:

$$\gamma(h) = \frac{1}{2|N(h)|} \sum_{i=1}^{|N(h)|} [Z(x_i) - Z(x_i+h)]^2 \quad (4.15)$$

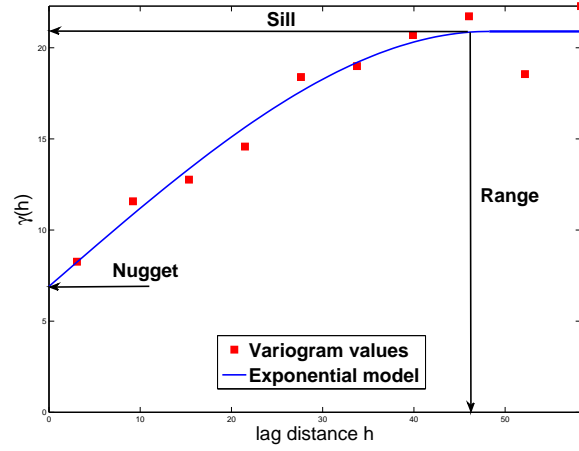


Figure 4.2: An example of an empirical semi-variogram values (in red) together with a best fitting (using least linear squares) exponential semi-variogram. The plot includes the locations of the semi-variogram nugget, sill and range.

where

- $N(h) \equiv \{(i, j) : x_i - x_j \leq h\}$
- $|N(h)|$ is the number of distinct elements of $N(h)$.
- $Z(x_i)$ is the measurement at position x_i .
- $Z(x_i + h)$ are all the measurements within a distance lag h

In Figure 4.2, the red square points are an example of an empirical semi-variogram plot defined by equation (4.15). The vertical axis is $\gamma(h)$ and the horizontal axis is a lag distance such that

$$\mathbf{h} = [3.0, 9.2, 15.3, 21.5, 27.6, 33.9, 39.9, 46.0, 52.2, 58.3]. \quad (4.16)$$

Each red square point represents an aggregation of neighborhood points where the quasi-stationarity assumptions hold. The measurements are acoustic modem Received Signal Strength (RSS) and will be discussed more thoroughly later in the chapter.

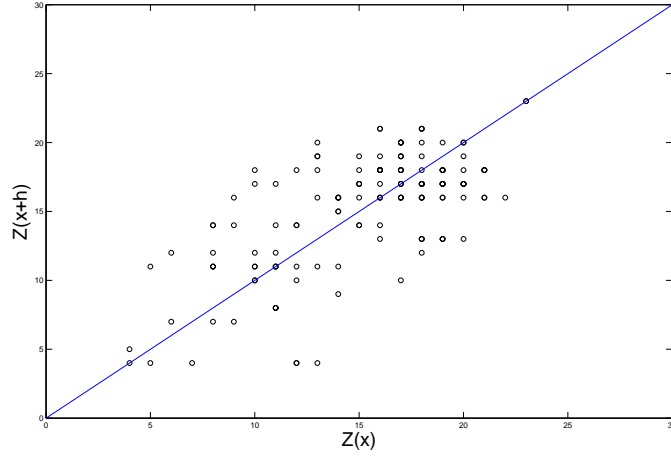


Figure 4.3: The h -scattergram of 152 RSS measurements taken from bearing window of 80 degrees. The lag distance (h) is 10 meters.

Comparison of covariance and variograms

Covariance is a measure of similarity between two random variables. Semi-variograms are a measure of dissimilarity. One way to see this is through an h -scattergram. It is a plot of the measurements $Z(x)$ versus the measurement plus a lag distance $Z(x+h)$. Figure 4.3 gives an example of an h -scattergram where the lag distance is equal to 10 meters.

The semivariance is the moment of inertia or spread of the h -scattergram about the 45-degree line. Equations for the covariance and correlation are:

$$\text{Covariance: } Cov(h) = \frac{1}{|N(h)|} \sum_{i=1}^{|N(h)|} [Z(x_i) \cdot Z(x_i + h) - \mu_o \cdot \mu_{+h}] \quad (4.17)$$

$$\mu_o = \frac{1}{|N(h)|} \sum_{i=1}^{|N(h)|} Z(x_i), \quad \mu_{+h} = \frac{1}{|N(h)|} \sum_{i=1}^{|N(h)|} Z(x_i + h) \quad (4.18)$$

$$\text{Correlation: } \rho(h) = \frac{Cov(h)}{\sqrt{\sigma_o \cdot \sigma_{+h}}} \quad (4.19)$$

$$\sigma_0 = \frac{1}{|N(h)|} \sum_{i=1}^{|N(h)|} Z(x_i), \quad \mu_{+h} = \frac{1}{|N(h)|} \sum_{i=1}^{|N(h)|} Z(x_i + h) \quad (4.20)$$

where μ_0 and μ_{+h} are the mean values for the tail and head values and σ_0 and σ_{+h} are the corresponding standard deviations. This is meant to highlight the relationship between the more traditional statistical measures and the semi-variance. Note that there is normally a maximum of the semi-variogram function $\lambda(\infty)$. As will be discussed in just a moment, it is called the sill. This is represented mathematically as:

$$\gamma(\infty) = \text{Var}\{Z(x)\} = C(0) \quad (4.21)$$

where $C(0)$ is the covariance between two points located at the same position. This will prove useful later on for the calculation of the kriging variance.

4.2.4 Variogram Data Modeling

Variogram data modeling is the process of selecting a function and determining the function parameters that properly fit the empirical semi-variogram data. Parametric and non-parametric functions are commonly used for this task.

Parametric Modeling

There are a number of parametric functions that are commonly used for summarizing the empirical semi-variogram. They include Linear, Spherical, Exponential, Gaussian and Matérn models. Figure 4.4 gives examples of these functions. The elasticity of a function is defined as the ratio of the percentage change in the function's output with respect to the relative change of its input. This can be represented by the following equation.

$$El[f(x)] = \frac{d \log(f(x))}{d \log(x)} \quad (4.22)$$

where $El[f(x)]$ is the elasticity of the function. These models have relatively little elasticity. This has positive and negative aspects.

On the positive side, the function's inflexibility is beneficial for a distance lag that is an

outlier—it can adequately capture the trend without being significantly impacted by a lag distance that may reflect several erroneous measurements collected in a single region. Also there tends to be a small number of parameters associated with these models. When using optimization techniques to determine parameter values, a model with fewer parameters is frequently easier and faster for calculating solutions.

On the negative side, if there is a lot of data for creating the variogram and there is a high level of confidence in the measurement process, it may be difficult to capture an accurate variogram with the limited elasticity of these parametric equations. As we will see the variogram directly impacts the interpolated estimate for the point of interest, the less accurate the variogram the less accurate the estimate becomes.

With regards to communication variograms, it is not usual to have data sets where signal strength are correlated at significantly different lag distances. For example, in Figure (4.4) the sine-hole function gives an example of this phenomenon. The evaluation of γ at the approximate lag-distances of 40 and 70 meters are the same ($\gamma(40) \propto \gamma(70)$). For this reason a non-parametric function may have greater utility than a parametric choice.

- Exponential

$$\gamma(h) = s[1 - e^{-h/r}] \quad (4.23)$$

- Spherical

$$\gamma(h) = \begin{cases} s[\frac{3h}{r} - \frac{1}{2}(\frac{h}{r})^3] & \text{if } h < r \\ s & \text{otherwise} \end{cases} \quad (4.24)$$

- Gaussian

$$\gamma(h) = s[1 - e^{-(h/r)^2}] \quad (4.25)$$

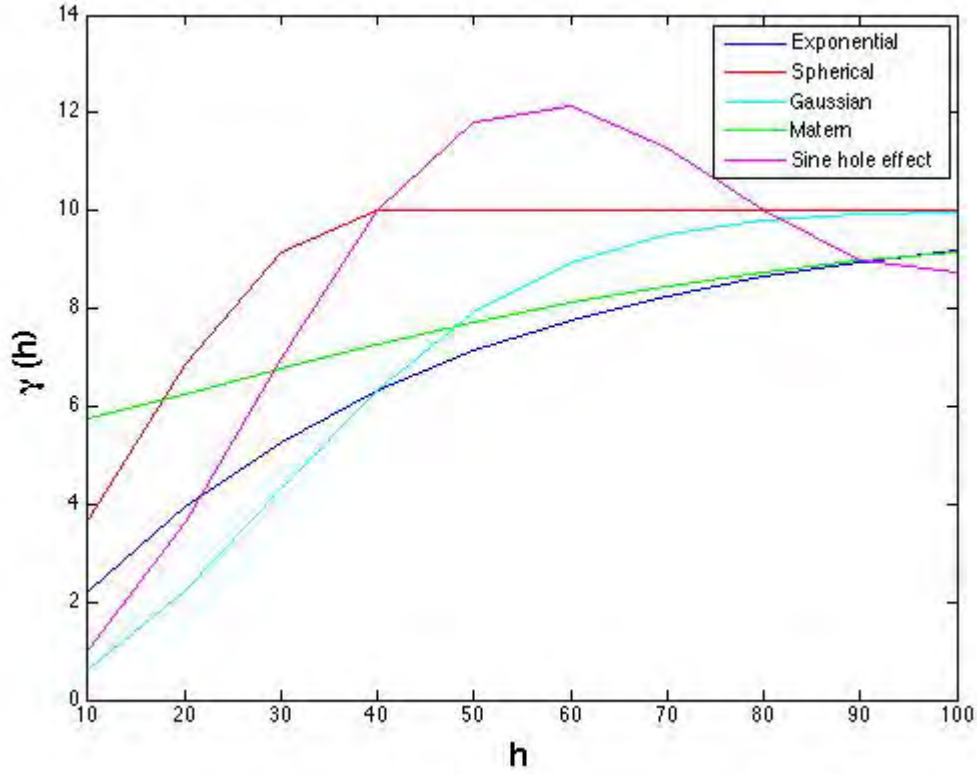


Figure 4.4: Examples of semi-variogram models.

- Matérn³

$$\rho(s; \xi, \nu) = \frac{1}{2^{\nu-1} \Gamma(\nu)} \left(\frac{2\nu^{1/2}s}{\xi} \right) H_{\nu} \left(\frac{2\nu^{1/2}s}{\xi} \right) \quad (4.26)$$

- Sine Hole Effect

$$\gamma(h) = s[1 - (r/h)\sin(h/r)] \quad (4.27)$$

Figure 4.2 shows an example of the fitting of the exponential model to the empirical semi-variogram. Least linear squares is used to determine the parameters of the exponential model.

³There are several different formulations of the Matérn function [20]. Later another version of the function will be used for representing covariance and signal strength functions.

Non-parametric Modeling

Typically, parametric models are used to ensure the conditionally negative definite property of the variogram. Still, there has been significant research into non-parametric modeling for variogram estimation. One approach by Shaprio and Botha [118] utilizes Bochner's Theorem [119] for developing conditionally negative definite variograms using quadratic programming. This technique has been extended by various researchers for modeling the covariance spectrum [120], [121], [122] and most recently by Huang et al. [23].

A recognized weakness of the traditional semi-variogram is the disproportional influence of data outliers. Another alternative approach for non-parametric estimation is to replace the lag-distance vector h with a more flexible nearest-neighbor parameter [21]. Yu et al. [22] highlight the difficulties with this approach (uniformly smooth changes, continuity over the range of distance data and a non-decreasing behavior) and introduce a variable nearest neighbor variogram estimator.

4.3 Conclusion

This chapter presented an introduction to structural analysis associated with spatial estimation. Of interest is a data-driven technique for interpolation. To accomplish this it is necessary to understand the spatial relationship between pairs of measurements. The variogram is a technique for understanding the dissimilarity between measurements over increasing distances. It was shown the relationship between variograms and covariance and why variograms are preferred with respect to spatial statistics.

The chapter included an empirical variogram with a least squares fitted exponential model. The next chapter leverages the structural analysis discussion by using the variogram function as part of the input into a spatial BLUE. This is known in the Geostatistical community as Kriging.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 5:

Communication Considerations for Variogram Estimation

The semi-variogram function summarizes the spatial dissimilarity between random variables. As will be seen, it will be used in a system of linear equations to determine an optimal estimation at a point of interest. The last chapter discussed how to construct a semi-variance function based on the spatial relationship between pairs of measurements. This chapter discusses the implementation of a semi-variogram model with respect to considerations for communications. It starts with looking at semi-variograms based on simulated data. Next, semi-variograms are built with respect to collected undersea acoustic modem data. Several parametric models are evaluated for fitting the empirical semi-variogram using a least linear square technique. These models are based on a lag-distances, they are range-centric.

The nature of the signal also permits consideration of a radial impact on estimation. In other words, unlike a gridded Euclidean distance metric, of primary interest is a source-based measurement model. There is an expectation of a peak measurement near the source location and it permits evaluation of two semi-variograms—one with respect to the range and a second with respect to the radial direction. The data used for estimating a radial semi-variogram comes from the values along an equidistant arc from the source. This chapter discusses the implications of a radial-centric semi-variogram and includes examples from acoustic modem datasets.

5.1 Simulated Semi-Variograms

Figure 5.1 is a simulated undersea acoustic modem dataset of received signal strength (RSS) measurements. The measurements were calculated using a Matérn function with parameters (25,25,10) at a center point (400, 400). The result is an isotropic field where the measurements degrade uniformly regardless of bearing (relative to the center point). Notice (unrealistically) that the set of measurements are complete. There are no gaps in coverage there is a measurement at each grid cell. The corresponding semi-variogram is given in Figure 5.2. It highlights several considerations for spatial estimation in communications.

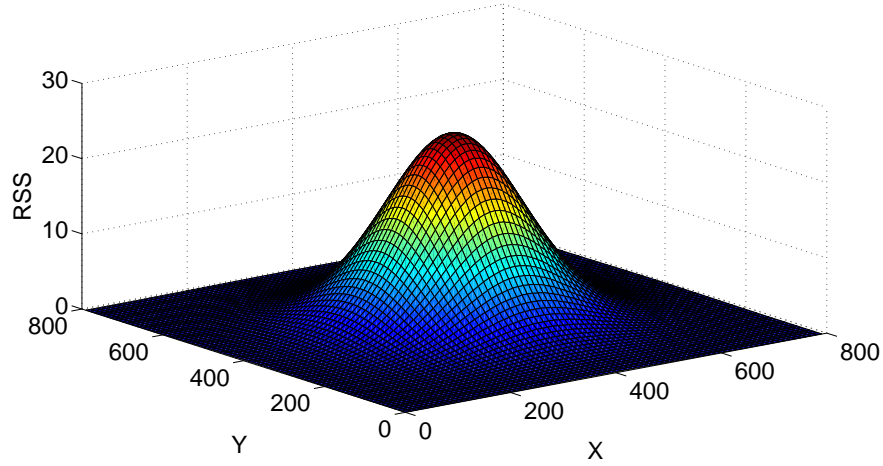


Figure 5.1: *Simulated acoustic modem RSS measurements (in dB).*

First, in an idealized scenario, there is no nugget. Recall that the nugget is the location in the plot where the semi-variogram crosses the y-axis. Intuitively, when the lag distance is zero ($h = 0$) the spatial dissimilarity should be zero ($\gamma(h = 0) = 0$). From a practical standpoint, there is frequently a nugget (a positive value at $\gamma(0)$) and it is assumed to represent measurement noise or error. Second, notice the sill is located near the maximum of the lag distance (the range is approximately 400). This can also happen with acoustic communication data, but frequently, there is a detectable sill with corresponding range that is commiserate with a quasi-stationarity assumption. In other words, there is frequently a plateau of the semi-variogram at 50-100 meters and the points that fall inside this range are the data points are normally selected for spatial interpolation.

5.2 Range Semi-Variograms

Of particular interest is the ability to communicate in confined areas. Figure 5.3 is an overhead photograph of Monterey Harbor, Monterey, CA. It is a confined and cluttered area where boat traffic, piers and moored boats can play havoc on acoustic signals. Figure 5.4 shows the results of undersea acoustic modem measurements collected in Monterey Harbor, Monterey, CA. The graphic highlights the static buoy located at position (N36.607024, W121.889747) and RSS measurements taken from a boat as it transited throughout the outer harbor. (Each data point in the graphic is the position of the boat and the corresponding RSS intensity using the MATLAB jet colormap.)

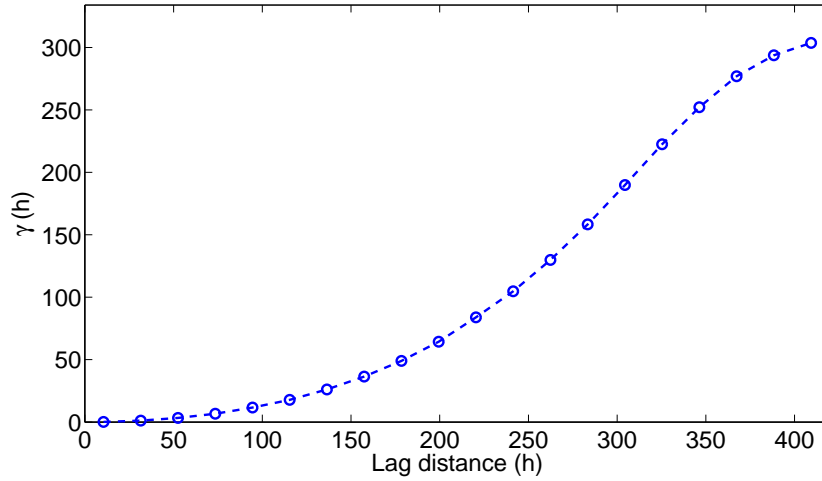


Figure 5.2: *The empirical semi-variogram calculated from the simulated acoustic modem RSS random field.*

Figures 5.5, 5.6, 5.7 and 5.8 show the empirical semi-variogram results looking at a sequence of 45-degree bearing windows. In each of the four cases one can notice the distinct differences in the plots which reinforces an anisotropic assumption. The top pair of plots in Figure 5.5 shows measurements that result in a traditional semi-variogram such that the range and sill are easily determined from the graph. The bottom pair of plots are also interesting. Even though there are fewer measurements than the other plots, one can clearly see that there is a drop in the semi-variogram at lag distances of approximately 100m. In other words, there was less dissimilarity in the data. This is not usual for signal propagation where for unknown reasons there is a drop in signal strength that recovers at greater lag-distances. If one has confidence in the measurements, the ability to match a function to the data may improve one's ability for prediction or interpolation. In curve fitting, one must balance the results of the empirical semi-variogram with a representative function that adheres to the restrictions of a conditionally negative definite form.

5.3 Curve Fitting the Empirical Range Semi-Variogram

Using the acoustic modem data set shown in Figure 5.4, Table 5.1 provides results from comparing four models (spherical, Matérn, Gaussian and bilinear) for semi-variograms built from seven bearing windows. The bearing window $([0,45])$ has less than 30 measurements and was not included. The models are evaluated by calculating the Root Mean



Figure 5.3: *An overhead picture of Monterey harbor where boat traffic, piers and moored boats can play havoc on acoustic signals.*

Square Error (RMSE) of the difference between the best fitting parameters (θ), using least squares for a candidate function ($f(h, \theta)$) and the semi-variogram values $\gamma(h)$. It is given by the following equation.

$$RMSE = \sqrt{\frac{\sum_{i=1}^n (f(h, \theta) - \gamma(h))^2}{n}} \quad (5.1)$$

The n is the number of the lag distance values (in this case $n = 7$). The results of the comparison show that the Gaussian model fit the different semi-variogram curves best but the bilinear and spherical models were not far behind. The worst performer was the Matérn model. It is also possible to improve the RMSE by looking for the best model of a collection. These results are shown in bold. The downside of this approach is the computational cost associated with the fitting the empirical semi-variogram to multiple models.

5.4 Radial Semi-Variograms

In this chapter and the last, there are examples of semi-variograms looking at an aggregation of data organized into lag-distances. The collected data was partitioned over several

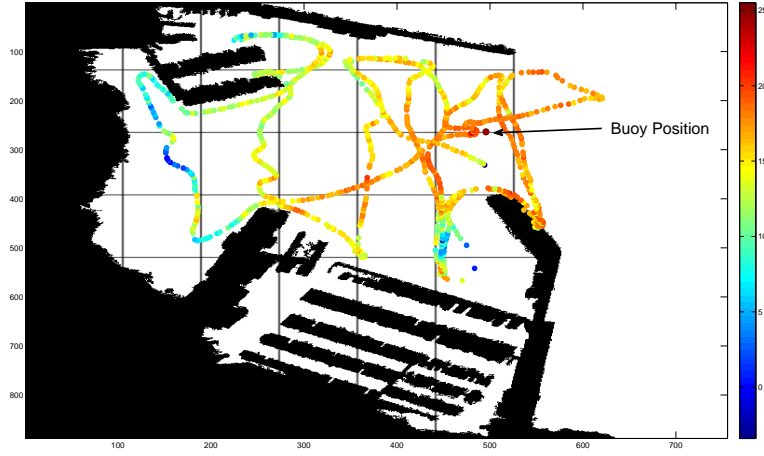


Figure 5.4: RSS measurements relative to an static buoy collected in Monterey harbor.

Model	Bearing	Spherical	Matérn	Gaussian	Bilinear
BW1	[45 90]	1.41	.63	1.01	1.34
BW2	[90 135]	3.27	5.64	3.56	4.25
BW3	[135 180]	5.21	6.29	3.55	4.07
BW4	[180 225]	1.22	2.92	.76	.86
BW5	[225 270]	.55	1.33	.59	.51
BW6	[270 315]	1.00	1.60	1.32	1.11
BW7	[315 360]	.58	1.00	.63	.57
Total		13.25	22.36	11.45	12.74

Table 5.1: The empirical semi-variograms curve fitting results of the four models. The models are evaluated using RMSE. The lowest error is highlighted in bold text.

bearing windows. The implicit assumption was that over each bearing window the region is homogeneous. This simplifying assumption produced a *range* semi-variogram.

In analogous fashion, it is possible to evaluate the relationship of the data relative to fixed radial distances from the center node. In an isotropic communications field, one would expect that the signal would remain uniform regardless of bearing. By aggregating data at fixed distances from the centroid, one can develop a semi-variogram function over increasing arc lengths. This will be called a *radial* semi-variogram.

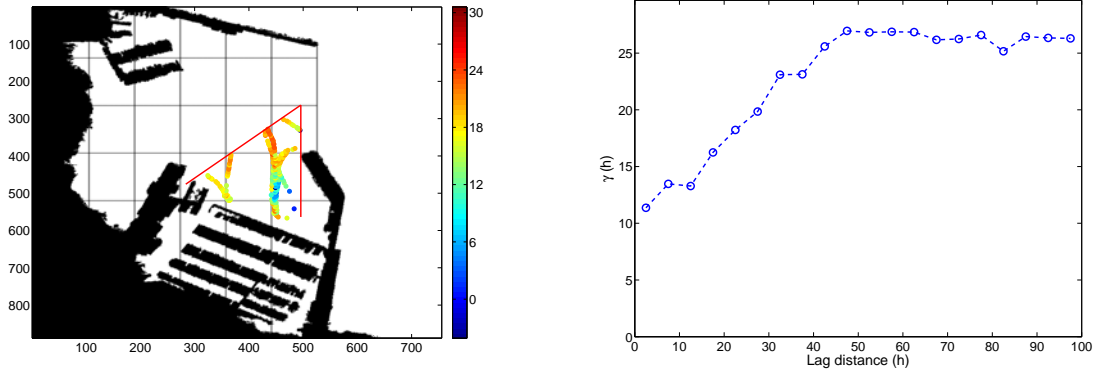


Figure 5.5: (Left) RSS measurements over a 45-degree bearing window. (Right) The resultant range semi-variogram. The maximum lag-distance for the semi-variogram is 100m. The nugget is approximately 11, the sill 26 and the range 50.

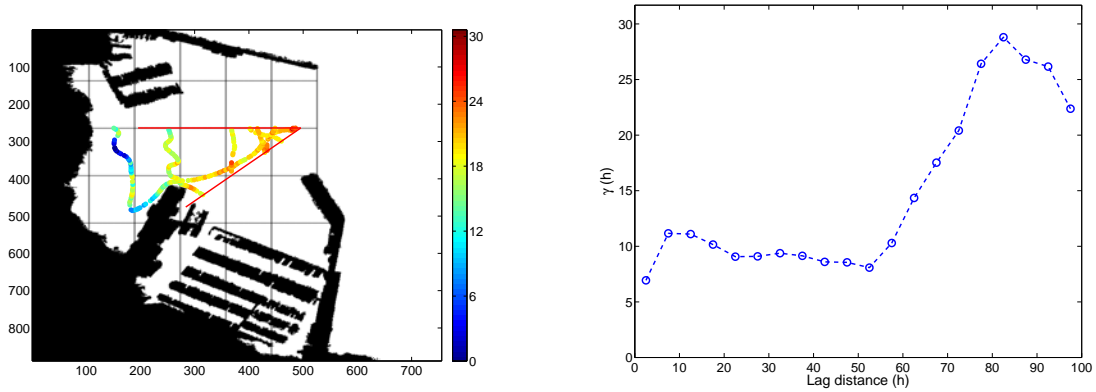


Figure 5.6: (Left) RSS measurements over a 45-degree bearing window. (Right) The resultant range semi-variogram. The maximum lag-distance for the semi-variogram is 100m. The nugget is approximately 6, the sill 26 and the range 80.

In other words, in the estimation process (that will be covered in the next chapter), if the measurements, within a radial ring (for example between 25 and 50m) from the center, had a tight uniform variance it would be useful to consider all these measurements for calculating the mean. The main advantage of this approach is that it leverages the nature of the signal. As a point energy source, it is possible to construct two distinct semi-variograms such that the combination of both semi-variograms may result in a better estimate through the use of a greater number of measurements.

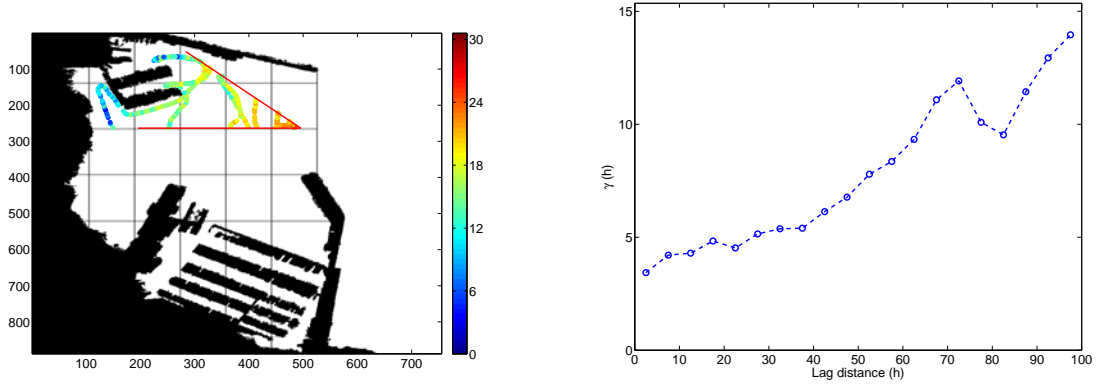


Figure 5.7: (Left) RSS measurements over a 45-degree bearing window. (Right) The resultant range semi-variogram. The maximum lag-distance for the semi-variogram is 100m. Selection of the sill and range is trickier here. The nugget is approximately 4, the sill is 13 and the range is 70, even though the semi-variogram increases after a lag distance of 80. The reason not to select a higher sill/range value is not to have a covariance function that influences the estimate with far away data points.

The general approach is to model the anisotropic behavior by partitioning the region. The bearing windows permit a limited ability to estimate the anisotropic random field, especially with respect to range. Similarly, the radial windows permit a limited ability to estimate the anisotropic random field, especially with respect to bearing. The limitations of the combined approach are associated with the selected dimensions of the bearing and radial windows, since it is assumed that over these windows the field is homogeneous.

In summary, radial semi-variogram technique effectively recognizes that measurements at equivalent radii but different bearings may be correlated. One might expect that the correlation might be particularly strong as the comparative distances from the centroid reduces to zero. The first step is to build the empirical radial semi-variogram.

5.5 Examples of Radial Semi-Variograms

The radial semi-variogram is constructed in the same way as the range semi-variogram with two exceptions. First, the radial semi-variogram is based on a fixed range from the center point. Second, normally the abscissa (element of an ordered pair plotted on the horizontal axis) is the lag distance that corresponds to euclidean distance measurements. In this case, the abscissa is increasing arc length. The possible range of values is $[0, \pi *$

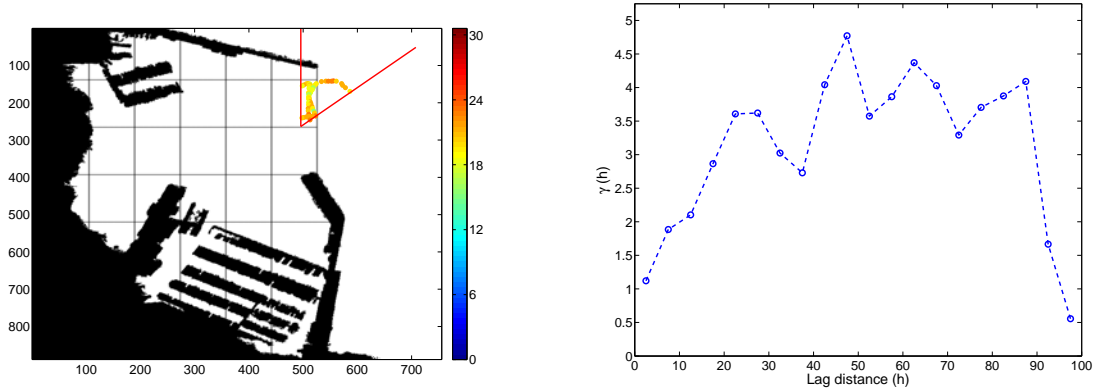


Figure 5.8: (Left) RSS measurements over a 45-degree bearing window. (Right) The resultant range semi-variogram. The maximum lag-distance for the semi-variogram is 100m. The nugget is approximately 1, the sill 5 and the range 50.

(circle diameter)]. This is normally discretized into a limited number of increasing arc lengths.

Figures 5.9-5.13 show the results for empirical radial semi-variograms based on the acoustic modem datasets S_1 - S_5 . Notice that each figure has six graphed empirical semi-variograms. These correspond to the discrete ranges (50, 100, 150, 200, 250, 300) such that measurements that are within $(-25, 25)$ of a particular range are used for building the radial semi-variogram. As anticipated, the radial empirical semi-variograms at ranges close to the centroid position tend to have lower semi-variance and they were consistent—as the radial bearing window increases the semi-variance did not noticeably change. This was due small dissimilarity between random variables within the radial ring.

At greater ranges, there was frequently a significant jump in the semi-variogram indicating greater variability with greater radial bearing lag distance. Of particular interest is the 200–250m plot in Figure 5.10. This shows an upward trend in the semi-variogram values as the bearing lag increases. This is consistent with characteristics of an anisotropic random field; as more data is considered from a larger arc length the variability increases. In this way, the radial empirical semi-variogram graphically illustrates a potentially anisotropic random field associated with communication channel estimation.

In summary, this chapter has described communication considerations for semi-variogram

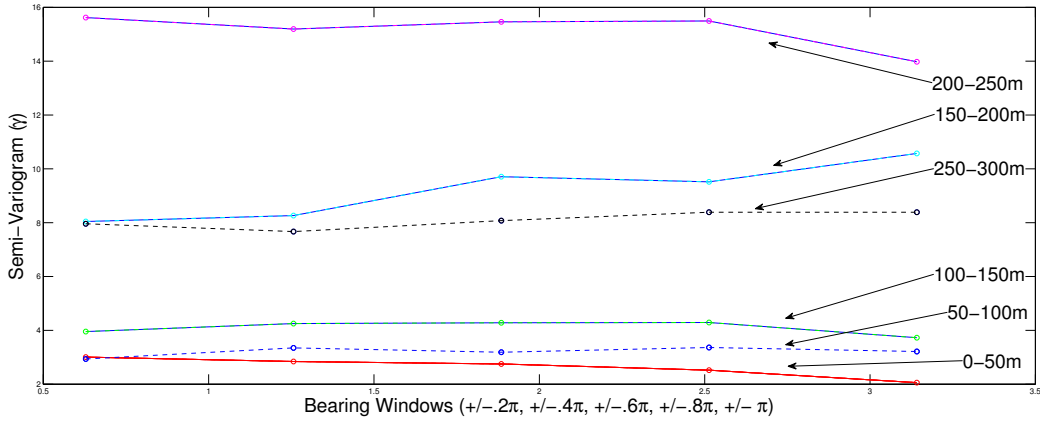


Figure 5.9: The radial semi-variograms from the S_1 acoustic modem data set collected in Monterey Harbor, July 2012. Each semi-variogram is labeled based on the fixed distance from the source position.

estimation. In particular the nature of the omnidirectional signal permits the data to be viewed in a polar perspective such that two semi-variograms are possible—one for range and a second for radial distance for the centroid. The combination of the two could potentially improve an estimation technique that is normally not available in a “traditional” Kriging application. In the next chapter, these techniques will be used in combination with a Best Linear Unbiased Estimator for producing stochastic estimates of points of interest.

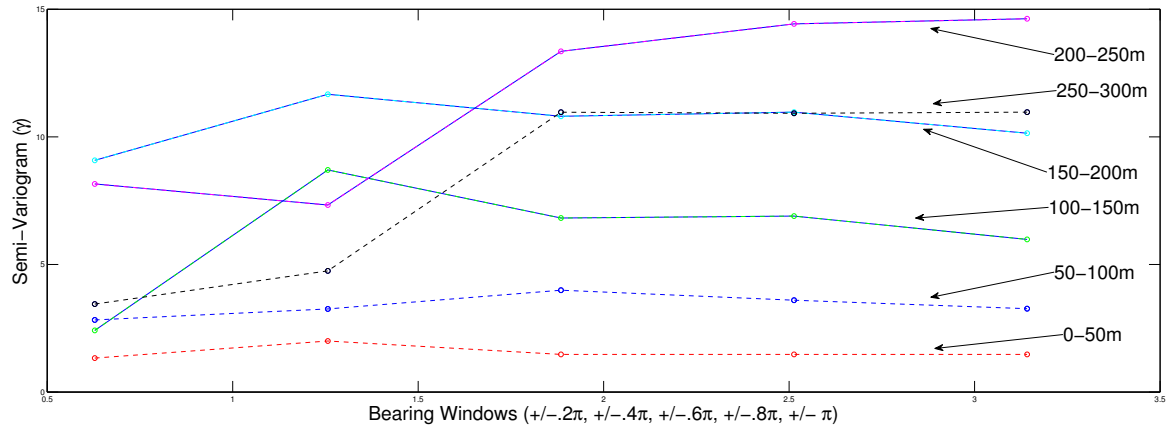


Figure 5.10: The radial semi-variograms from the S_2 acoustic modem data set collected in Monterey Harbor, July 2012. Each semi-variogram is labeled based on the fixed distance from the source position.

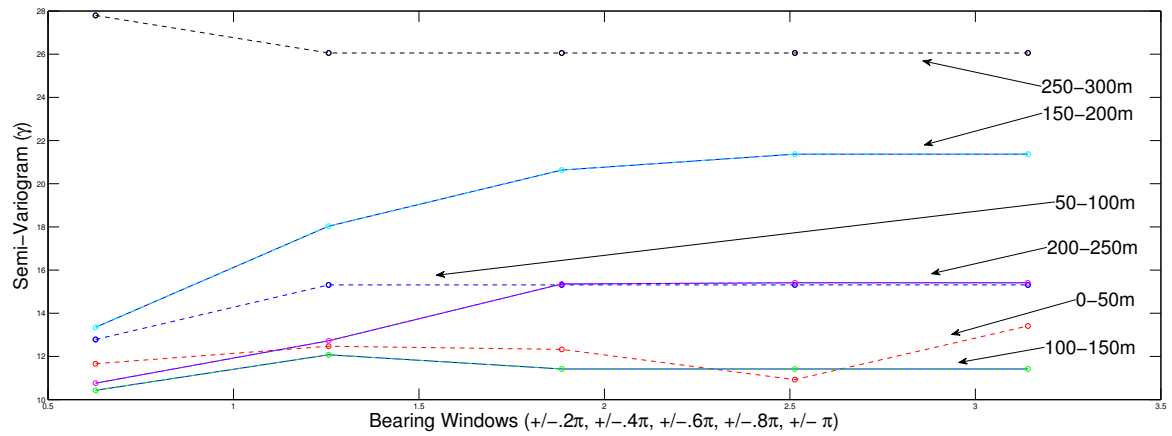


Figure 5.11: The radial semi-variograms from the S_3 acoustic modem data set collected in Monterey Harbor, July 2012. Each semi-variogram is labeled based on the fixed distance from the source position.

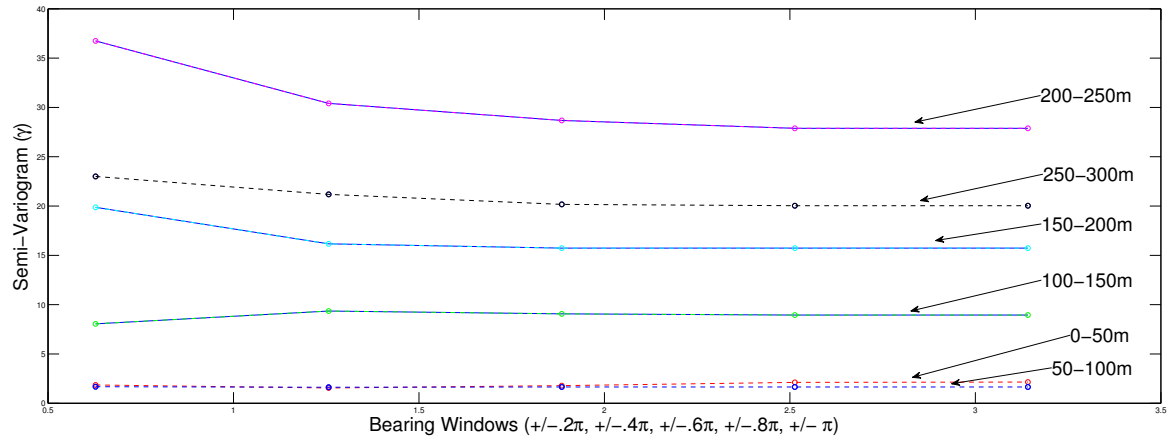


Figure 5.12: The radial semi-variograms created from the S_4 acoustic modem data set collected in Monterey Harbor, July 2012. Each semi-variogram is labeled based on the fixed distance from the source position.

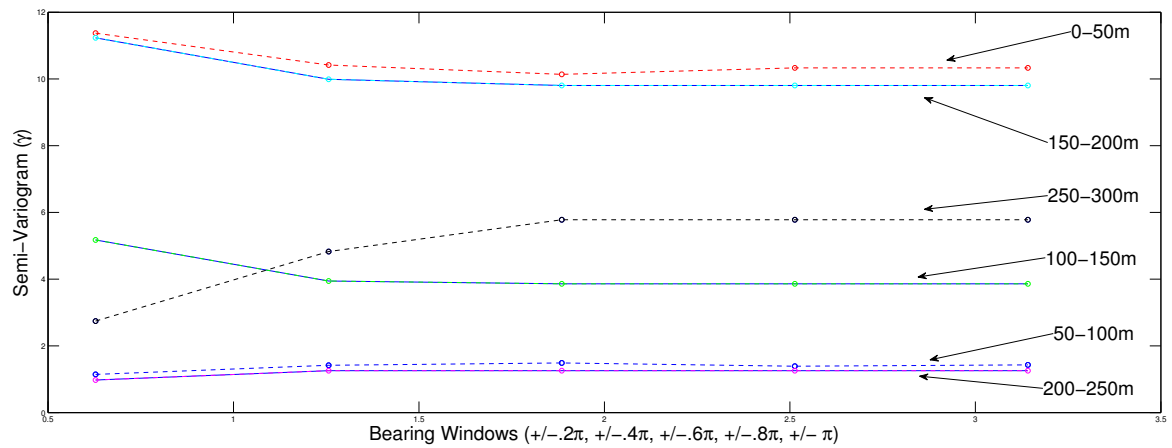


Figure 5.13: The radial semi-variograms created from the S_5 acoustic modem data set taken from Monterey Bay in July 2012. Each semi-variogram is labeled based on the fixed distance from the cell center.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 6:

Communications Kriging

In previous chapters, variogram analysis was introduced for a data-driven methodology for calculating a semi-variogram. This chapter will use the results of prior chapters and introduce an interpolation method for calculating an estimate for a point of interest (POI). When the POIs cover an entire bounded domain, the result is a stochastic estimate for a random field. The methodology is called Kriging. It is particularly useful in the building of signal strength or connectivity maps since it produces both a mean and variance estimate.

The chapter starts with an overview on Kriging including the derivation of the ordinary Kriging equation. It is followed by the definition of the Local Connectivity Map (LCM) and a discussion of communication considerations with respect to pertinent signal statistics. The chapter concludes with an introductory section of undersea acoustic communications and how it differs from above-sea wireless networks.

6.1 Kriging

The term Kriging was coined by Georges Matheron [114] to recognize initial contributions made by South African mining engineer Daniel Krige [16]. It uses measurements at known locations to produce spatially interpolated mean and variance prediction at locations of interest. This is accomplished through the variogram model that produces distance based correlations with other data measurements.

The ordinary Kriging estimator can be derived with Optimal Estimation Theory [123]. The model of a spatial random variable (Z) is a linear combination of an observed set of values:

$$\text{Model Assumption: } Z(x) = \mu + \varepsilon(x) \quad (6.1)$$

where μ is a mean value and ε is a zero mean error with known covariance structure $Cov(Z(x_i), Z(x_j)) = K_\theta(x_i, x_j)$, and θ is a vector of spatial covariance parameters.

Kriging is called a Best Linear Unbiased Estimator (BLUE). It is linear since the predictor takes the form:

$$\text{Predictor Assumption: } \hat{Z}(x_0) = \sum_{i=1}^n \lambda_i Z(x_i) \text{ Subject to: } \sum_{i=1}^n \lambda_i = 1, \lambda_i \geq 0 \quad (6.2)$$

The intuition is that there are several sensor measurements available at known locations and one is interested in producing an estimate (\hat{Z}) at a nearby point of interest (x_0). To produce an estimate, one may assign fractional values to each of the measurements that take into consideration the spatial distances to the location of interest and a correlation function. Note that the predictor constraint ensures that the fractional values add to one.

The goal is to determine the values of λ_i that ensure the variance is minimized and the estimator is unbiased. Using the quasi-stationarity assumption, the expected value of the difference between any random variables in a defined neighborhood is zero.

$$E[Z(x_i) - Z(x_j)] = 0 \quad (6.3)$$

The condition for which the estimate is unbiased is $\sum_{i=1}^n \lambda_i = 1$. This constraint ensures uniform unbiasedness since

$$E[\hat{Z}(x_0)] = \sum_{i=1}^n \lambda_i E[Z(x_i)] = \sum_{i=1}^n \lambda_i \mu = \mu \sum_{i=1}^n \lambda_i = \mu \quad (6.4)$$

To construct an optimal estimator for $Z(x_0)$, one can minimize the mean square prediction error (MSE). This is calculated by finding the expected value ($E[\cdot]$) of the difference squared from the true value ($Z(x)$) and the estimate ($Z(x_0)$).

$$MSE = E[Z(x) - Z(x_0)]^2 = E[Z^2] - 2E[Z\hat{Z}] + E[\hat{Z}^2] \quad (6.5)$$

To simplify the notation, let $Z(x) = Z$ and $\hat{Z}(x_0) = \hat{Z}$. Expanding the three right-hand side terms and simplifying the prediction assumption (6.2) to $\hat{Z} = \lambda^T \mathbf{Z}$ yields:

$$E[\hat{Z}^2] = E[(\lambda^T \mathbf{Z})^2] = E[\lambda^T \mathbf{Z} \mathbf{Z}^T \lambda] = \lambda^T E[\mathbf{Z} \mathbf{Z}^T] \lambda = \lambda^T (\mathbf{K} + \mu^2 \mathbf{I}) \lambda \quad (6.6)$$

$$E[\hat{Z} \mathbf{Z}] = E[\lambda^T \mathbf{Z} \mathbf{Z}] = \lambda^T E[\mathbf{Z} \mathbf{Z}] = \lambda^T (\mathbf{k} + \mu^2 \mathbf{1}_n) \quad (6.7)$$

$$E[Z^2] = \mu^2 + 2\mu\varepsilon + \varepsilon^2 \quad (6.8)$$

where \mathbf{K} is a symmetric covariance matrix of the selected data points and \mathbf{k} is a column matrix of covariances between the POI and the selected data points and $\mathbf{1}_n$ is a $n \times 1$ column vector of ones. The combination of these terms together with the Lagrange multiplier constraint, yields the following equation:

$$MSE = \lambda^T (\mathbf{K} + \mu^2 \mathbf{I}) \lambda + \lambda^T (\mathbf{k} + \mu^2 \mathbf{1}_n) + \mu^2 + 2\mu\varepsilon + \varepsilon^2 + 2(\lambda^T - \mathbf{1})v \quad (6.9)$$

where v is the Lagrangian multiplier. The minimization is achieved by taking the derivative of the MSE with respect to λ and setting it equal to zero and solving ⁴:

$$\frac{1}{2} \frac{\partial MSE}{\partial \lambda} = [\mathbf{K} + \mu^2 \mathbf{I}] \lambda - \mathbf{k} - \mu^2 \mathbf{1}_n + v \mathbf{1} = 0 \quad (6.10)$$

$$\mathbf{K} \lambda - \mathbf{k} + v \mathbf{1} = 0 \quad (6.11)$$

$$\mathbf{K} \lambda + v \mathbf{1} = \mathbf{k} \quad (6.12)$$

By incorporating the unbiased constraint within the matrices \mathbf{K} and \mathbf{k} , it yields the follow-

⁴The first right-hand side term of (6.10) leverages the matrix identity $\frac{\partial}{\partial \mathbf{x}} (\mathbf{x}^T \mathbf{W} \mathbf{x}) = 2\mathbf{W} \mathbf{x}$

ing matrix equation

$$\mathbf{K}\boldsymbol{\lambda} = \mathbf{k} \quad (6.13)$$

Since the covariance matrix is assumed to be positive definite, there exists an inverse and ensures a solution for $\boldsymbol{\lambda}$. The derivation of the linear equation emphasizes the covariance matrix. This is consistent with similar optimal estimation techniques (such as the Kalman Filter). To use the Kriging Equation (6.13) it is necessary to “convert” the variogram evaluation at a designated lag distance $\gamma(h)$ to a covariance value $C(h)$. This is accomplished using the relationship between the covariance and variogram developed in the previous chapter, $C(h) = C(0) - \gamma(h)$. Note that the same Equation (6.13) can be derived using the variogram [11].

In summary, with given measurements $\{z(x_1), z(x_2), \dots, z(x_n)\}$ of the spatial random field Z at locations x_1, x_2, \dots, x_n , the optimal ordinary Kriging predictor at point x_0 is $\hat{Z}(x_0) = \sum_{i=1}^n \lambda_i Z(x_i)$. To determine the values for $\boldsymbol{\lambda}$ requires solving a system of linear equations⁵:

$$\begin{bmatrix} C(s(x_1, x_1)) & \dots & C(s(x_1, x_n)) \\ \vdots & & \vdots \\ C(s(x_n, x_1)) & \dots & C(s(x_n, x_n)) \\ 1 & \dots & 1 \end{bmatrix} \begin{bmatrix} \lambda_1 \\ \vdots \\ \lambda_n \\ \nu \end{bmatrix} = \begin{bmatrix} c(s(x_1, x_0)) \\ \vdots \\ c(s(x_n, x_0)) \\ 1 \end{bmatrix}$$

The derivation for the variance for the ordinary Kriging model can be found in [124] and is given by:

$$\sigma^2 = C(0) - \sum_{i=1}^n \lambda_i C(x_i - x_0) - \nu \quad (6.14)$$

⁵An entry in the matrix \mathbf{K} is $C(h)$, where C is a covariance function and h is a lag distance. A similar argument holds for the elements of the column matrix \mathbf{k}

The calculation of the covariance function is based on a distance function $s(\cdot, \cdot)$ between points x_i and x_j for entries in \mathbf{K} and between points x_0 and x_j for entries in \mathbf{k} . The distance function will be discussed shortly.

There are many variants of Kriging. Three of the most common types of Kriging include: simple, ordinary and universal. Simple Kriging assumes that the mean function is zero ($\mu(x) = 0$). Ordinary Kriging (derived above) is when the mean process is assumed to be an unknown constant ($\mu(x) = c$) and universal Kriging describes the case of a more general mean function (also called Kriging with a trend).

6.1.1 Data transformation

For most problems in geostatistics, a Euclidean distance metric is used for determining the spatial relationship between two points since measurements are typically made on a cartesian coordinate plane. For building a LCM, there is a slightly modified approach. This can be illustrated through two simple examples.

In an open, flat environment, signal strength may degrade uniformly over increasing distance relative to a fixed central reference point. The data is isotropic and two measurements may be highly correlated even if the Euclidean distance showed the points were far apart. In a cluttered environment, the signal strength can degrade non-uniformly over increasing distance. In other words, the data is anisotropic. For this reason, any comparison between measurements should take into consideration the difference of the distances to the source node. This is given by:

$$s(x_i, x_j) = ||x_i - x_c|| - ||x_j - x_c|| \quad (6.15)$$

where x_c is the source node.

6.1.2 An Example

At this point it is instructive to build a Kriged model of simulated signal statistics. Figure 6.1 shows a plot of simulated received signal strength (RSS) measurements. Both the x and y axis are in meters. The source of the transmissions is located at position (1000,1000).

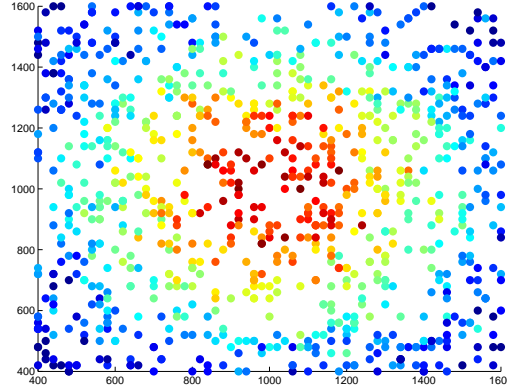


Figure 6.1: *The plot represents a simulated data set of Received Signal Strength (RSS) acoustic modem measurements. The measurements are taken relative to a centrally located static buoy as a source node. The measurement intensity values range from 1 to 30dB using the MATLAB colormap jet.*

The RSS intensities are based on the MATLAB colormap jet template. The signal measurements have a range of approximately $[1, 30]$ such that top values are red with a maximum of 30dB and transitions to orange, yellow and light blue as RSS values decrease. Dark blue represents the low end of the spectrum.

Figure 6.2 shows an example of an iteration in the procedure. Suppose one is interested in an estimate of the signal strength at a POI. For this example the POI is (840,600). The first step is to build a semi-variogram. It is assumed that the signal strength does not degrade uniformly over bearing and therefore the data is anisotropic. A bearing window is selected around the POI for building the semi-variogram. This is shown by the red lines in Figure 6.2a. The resulting empirical semi-variogram is in Figure 6.2b. It has been fitted using a least square process minimizing errors between a Gaussian function (in blue) and sum of the mean squared difference of the lag vector (red squares).

The next step is to identify a neighborhood of points around the POI to use for the estimate. There are many strategies for doing this. One goal is to minimize any “screening effect”, this happens when points selected close to the POI gain a majority of the Kriging weights. These points screen the potential influence of others used in the estimate. This suggests an overall procedure (that is consistent with least square theory) to select sample points that are well-dispersed.

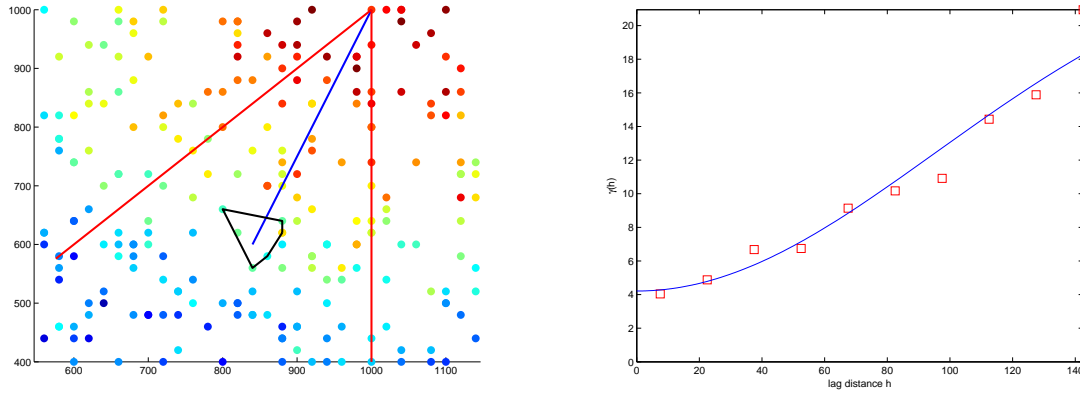


Figure 6.2: Figure 6.2a shows the collection of points used to create the variogram. Figure 6.2b shows the resulting semi-variogram.

Strategies for neighborhood selection can be generalized to methods seeking to approximate or simplify the Kriging equations or the selection of a flexible set of covariance functions. These techniques have been recently developed to address spatial interpolation of large datasets. Covariance tapering [125] is an example of the former, and the goal is to taper the covariance to zero beyond a defined range. This permits the systems of equations to be solved more efficiently. Fixed rank Kriging [126] is an example of the latter. It relies on a spatial random effects model such that the process can be expressed as a linear combination of spatial basis functions.

In the example, a simple expanding ring search approach was taken. The expanding ring of fixed distances selects dispersed points that ensure adequate angular and distance coverage. In this case, a neighborhood of six points was obtained. The convex hull around the points are given in Figure 6.2a. Table (6.1) shows the (x,y) locations of the points and the RSS values. The next step is to calculate the Kriging mean and variance. As described earlier, it is necessary to determine covariance values from the variogram function using the relation between the two (4.10) presented in the chapter on structure analysis.

x_1	x_2	z
880	620	20.07
820	620	15.04
840	560	14.39
860	580	12.43
880	640	14.39
800	660	13.59

Table 6.1: A table of six points consisting of position and RSS measurement used for the Kriging mean and variance example.

The written-out full system of equations for this example is as follows:

$$\begin{bmatrix} \lambda_1 \\ \lambda_2 \\ \lambda_3 \\ \lambda_4 \\ \lambda_5 \\ \lambda_6 \\ \lambda_7 \end{bmatrix} = \begin{bmatrix} 15.6743 & 15.2465 & 11.8903 & 14.0170 & 15.3526 & 15.6597 & 1 \\ 15.2465 & 15.6743 & 13.7632 & 15.2363 & 14.2379 & 15.0785 & 1 \\ 11.8903 & 13.7632 & 15.6743 & 15.1027 & 10.0376 & 11.5084 & 1 \\ 14.0170 & 15.2363 & 15.1027 & 15.6743 & 12.4805 & 13.7223 & 1 \\ 15.3526 & 14.2379 & 10.0376 & 12.4805 & 15.6743 & 15.4738 & 1 \\ 15.6597 & 15.0785 & 11.5084 & 13.7223 & 15.4738 & 15.6743 & 1 \\ 1 & 1 & 1 & 1 & 1 & 1 & 0 \end{bmatrix}^{-1} \begin{bmatrix} 14.7650 \\ 15.5786 \\ 14.4707 \\ 15.5475 \\ 13.4848 \\ 14.5332 \\ 1 \end{bmatrix} = \begin{bmatrix} 0.0001 \\ 0.5334 \\ 0.0000 \\ 0.4727 \\ 0.0000 \\ 0.0001 \\ 0.0000 \end{bmatrix}$$

The derived solution for the system of equations has no constraint such that the sum of the values for lambda can include negative numbers. While there are different opinions, it is the author's opinion that a negative weight does not make intuitive sense. The minimum impact a measurement should have on the estimate is zero. This is consistent with the 1984 paper by Barnes and Johnson [127]. Since the sum of the weights is one, the impact of having nonnegative values is that the Kriging weights are all on the interval [0,1]. A second modification for solving the system of equations is to use a Non-linear Least Squares (NLS) technique, where the initial guess for the solution is the normalized inverse distance of all the measurements from the POI. For this particular example, the residual of the NLS solution is 0.00057. From Equations (6.2) and (6.14), the Kriging mean and variance solution is:

$$\sum_{i=1}^n \lambda_i Z(\mathbf{x}_i) = \begin{bmatrix} 0.0001 \\ 0.5334 \\ 0.0000 \\ 0.4727 \\ 0.0000 \\ 0.0001 \end{bmatrix}^T \begin{bmatrix} 20.07 \\ 15.04 \\ 14.39 \\ 12.43 \\ 14.39 \\ 13.59 \end{bmatrix} = 13.90 \quad (6.16)$$

$$K(0) - \sum_{i=1}^n \lambda_i K(\mathbf{x}_i - \mathbf{x}) - v = 18.00 - \begin{bmatrix} 0.0001 \\ 0.5334 \\ 0.0000 \\ 0.4727 \\ 0.0000 \\ 0.0001 \end{bmatrix}^T \begin{bmatrix} 14.76 \\ 15.57 \\ 14.47 \\ 15.54 \\ 13.48 \\ 14.53 \end{bmatrix} - 0 = 2.3463 \quad (6.17)$$

Figure 6.3 shows the results from applying the procedure to all POIs which, in this case, are points every ten meters along the grid from [400,400] to [1600,1600].

6.1.3 Convergence

An appropriate question to ask involves the convergence and rate of convergence of the Kriging technique. That is, as the number of measurements in the domain increase, does the Kriging estimate of POIs approach the “true” value and at what rate does this convergence occur? Intuitively, it makes sense that as the number of measurements moves to infinity, the probability approaches one that there will be a measurement at the same location as the POI. In the set of linear equations, the λ_i weight for the measurement would be one while all others would be zero and this would be a valid solution with zero error.

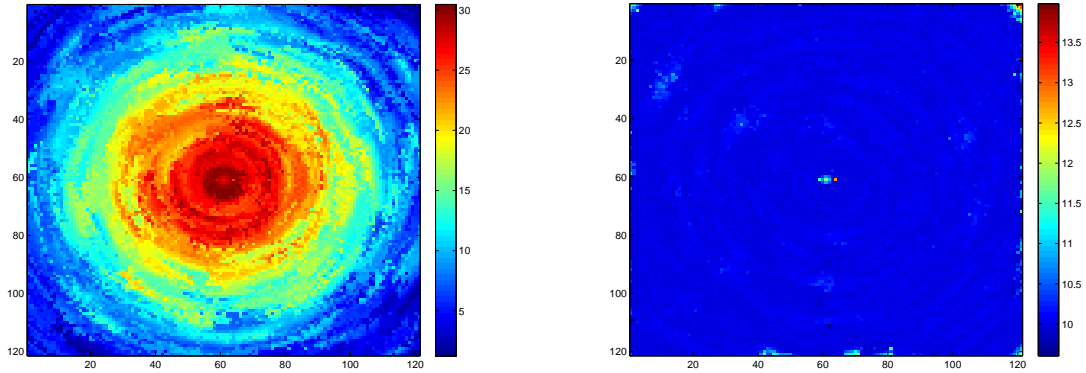


Figure 6.3: *The figures represent the Kriging mean and variance (on the left and right, respectively). The original input data is shown in Figure 6.1.*

Stein has published extensively on asymptotic properties of linear predictors [128], [129]. An emphasis in his approach is a comparison of a correct and incorrect second order structure (a second order Gaussian structure is defined by the mean and covariance pair) and looking at the spectral representation of the prediction errors. He points out that it is possible to have a poor covariance function and still get an accurate prediction. This is exemplified by the situation described in the prior paragraph—a dense number of measurements near the POI. A particularly insightful point is that, if this is the case, the low spectrum frequency has little impact on the prediction, it is the higher spectrum that is more influential. This is the case for interpolation but not for extrapolation. Since the focus is on the dissertation is interpolation, “...the focus in model selection and estimation of spectra should be on the high frequency behavior.” [129]

Second, Stein discusses the differences in asymptotic analysis between the communities interested in times series analysis and spatial analysis. With the former, the natural approach is to increase the observation region as the number of observations grow. This permits the distance between observations to stay roughly the same. He calls this increasing-domain asymptotics. For spatial analysis, an alternate approach that may be of more interest is the analysis associated with increasing observations over a bounded region. He calls this fixed-domain asymptotics.

Staying with fixed-domain asymptotics, Table 6.2 shows Kriging error results where a

Number of points	Point Density (%)	L_1 Mean Error
4	.00318	0.3983
8	.00637	0.1162
16	.0127	0.0642
32	.0254	0.0505
64	.0509	0.345
96	.0763	0.0286
128	.1019	0.0247
256	.2037	0.0152

Table 6.2: *The table shows the results of a Monte Carlo simulation of 1000 runs estimating the error based on increasing number of points. Given an assumed true model and true semi-variogram, it shows the convergence of the Kriging technique as the number of points in a fixed domain increases.*

“true” dataset was created using a Matérn function with a center point in a bounded square region 800 meters square. The dataset consists of a point every meter squared. A second dataset was created which centered on position (750, 750) where all points fit within a disk with radius 200. This data set had added random noise ($N(0, \sigma^2 = 6.25)$). A Kriging estimate was calculated using a neighborhood of eight points or less. The error was calculated by looking at the absolute value of the difference between the Kriging estimate and the true value (the L_1 norm). This was done for several iterations with an increasing density of points (4,8 16,32,64,96,128,256). This process was repeated 1000 times as a Monte Carlo simulation. Figure 6.4a shows the resulting plot.

The results show the convergence of the Kriging estimator as the number of points available for an estimate increases. This includes an assumption that the covariance function is perfectly known. The mean value for each group is highlighted by a red circle (Figure 6.4b) and is given in the third column of Table 6.2. By connecting the points, one is able to estimate the rate of convergence. This is shown in Figure 6.4b. Notice that the Kriging errors are unrealistically small. This is due to principally two reasons: First semi-variogram is perfectly accurate and there is no measurement error (there is no nugget) represented in the semi-variogram model. Second, the algorithm for finding the neighborhood of points attempts to select points uniformly around the POI. It tends to create a good estimate, since it tends to average out any trend to the data—as is the case with signal strength data sets.

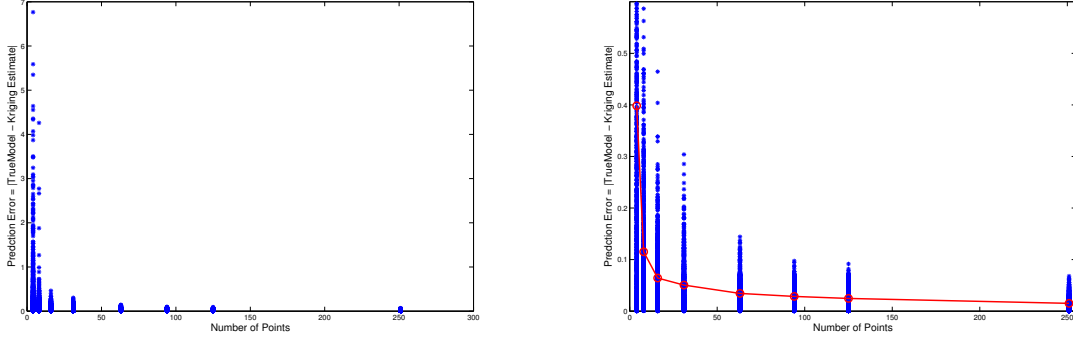


Figure 6.4: The plots show the results of a Monte Carlo simulation calculating the Kriging error that is the difference between a true model and a model with added random noise. The red line on the right hand side shows the mean value for each bin and the line gives an indication of the rate of convergence.

6.1.4 Stochastic Averaging of Kriging Models

As described in the previous chapter, it is possible to build two semi-variograms that are relevant to Kriged communication mapping, a *range* variogram and a *radial* variogram. Each makes simplifying assumptions about aggregating data. In the case of the radial semi-variogram, a limited number of bearing windows are defined over the interval $[0, 2\pi]$. While for the range semi-variogram the aggregation happens over a discrete vector of ranges. It is based on using the source node x_c as a reference point such that:

$$\begin{bmatrix} r \\ \psi \end{bmatrix} = \begin{bmatrix} ||x - x_c|| \\ \text{atan2}(x, x_c) \end{bmatrix}$$

The question is how to combine the measurements into a Kriging model and when it is appropriate. There are a couple of ways of doing this. The first is to construct a mean function which has range and bearing as inputs (e.g., $f(r_\psi, \theta_r)$). The difficulty is that we are particularly interested in the potential anisotropic behavior of the random field and using a function with these inputs causes one of two undesirable results. First, if all data is used, it produces a “watered down” mean and variance estimate that has the effect of averaging out angular anisotropic behavior.

Second, if range and bearing data is used from a bearing window it limits the impact of the bearing data. For example, if data was aggregated from a 50m full circular band (+/- 25m) and it produced a mean with a tight variance it would be useful to incorporate this into the overall estimate for each bearing window at 50m. An alternative methodology is a linear combination the measurements. Consider combining together random processes that can be written as a summation of functions such that

$$F(x) = \sum_{i=1}^p \alpha_i f_i(x_i) \quad (6.18)$$

where x_i is the i^{th} component of a d-dimensional vector x and f_i is an arbitrary univariate function and α_i is a scalar or gain.

In this case, the random process is a sum of random functions

$$Z(x) = \beta_1 Z(x_r) + \beta_2 Z(x_\theta) + \varepsilon_r + \varepsilon_\theta \quad (6.19)$$

where the random field $Z(x_r)$ corresponds to the range estimate and $Z(x_\theta)$ corresponds to bearing and $\beta_1 + \beta_2 = 1$. The linear combination of the means can be scaled by the appropriate variance as follows:

$$\hat{Z}(x) = \frac{\sigma_\theta^2}{(\sigma_\theta^2 + \sigma_r^2)} f(x_r) + \frac{\sigma_r^2}{(\sigma_\theta^2 + \sigma_r^2)} f(x_\theta) \quad (6.20)$$

and the minimum mean square estimation error is given by:

$$\sigma^2 = \left(\frac{1}{\sigma_\theta^2} + \frac{1}{\sigma_r^2} \right)^{-1} \quad (6.21)$$

This calculation is done for each POI.

Figure 6.5, shows a random field estimate using the linear combination of range and bearing regression models with the same simulated acoustic modem data that was used to produce Figure 3. The range covariance uses the Matérn function and the radial function use the semi-variogram results from figure (10) converted into a covariance function.

This stochastic averaging of Kriging models is similar to a technique known as Additive Kriging. This was recently discussed in a paper by Durrande et al. [130]. It builds on modeling techniques collectively known as Generalized Additive Models (GAM) [131]. The additive Kriging equations are as follows:

$$\mu(x) = \mathbf{k}(x)^T (\mathbf{K}_1 + \mathbf{K}_2)^{-1} \mathbf{z} \quad (6.22)$$

$$\sigma^2 = \mathbf{K}_1(x, x) - \mathbf{k}(x)^T (\mathbf{K}_1 + \mathbf{K}_2)^{-1} \mathbf{k}(x) \quad (6.23)$$

As can be seen these are additive changes for ordinary Kriging equations. The technique is specifically for multiple regression problems where the mean and variance can be decomposed into the equal sum of parts (although Durrande suggests that additive Kriging models are useful even when the trend is not purely additive). Conceptually, the above described technique of combining estimates of a single regression problem is very similar to Additive Kriging, But is different for two reasons. First it is a single linear regression problem over a single data set. Second, the means and variance are not additive without variance-derived gains since the results would overestimate the moments.

By subtracting the variance matrix estimates and summing along both axes of the difference matrix (summing along both axes is called the integral image), one can compare the combinatorial range and radial approach versus the range only estimate. The overall improvement was 33.9% in terms of a reduction in overall variance. One note of caution, the simulated dataset is isotropic and this produces a nearly linear radial covariance functions such that, if necessary, a point (from the neighborhood selection process) can be far

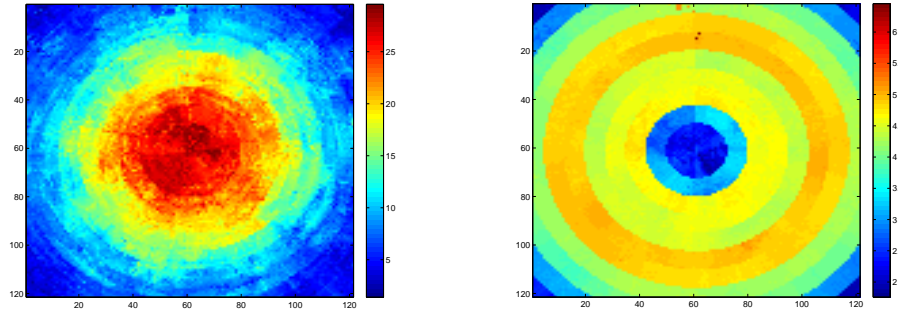


Figure 6.5: *Mean and variance estimates of the RSS random field using the additive Kriging methodology.*

from the POI and contribute towards a low variance estimate in the Kriging solution. Still under many situations the use of this technique may significantly improve/reduce model uncertainty.

Two additional points: First, this is easily extended into three-dimensional modeling where the “ring” lag distance is replaced by a torus of designated thickness. Second, this is a technique that is unique to a point source sensor model. In other words, it doesn’t have utility in the x,y,z euclidean metric space. In this sense, it is somewhat specific to point source events.

In summary, when the data is available, it is useful to combine a radial and range Kriging estimates through a stochastic additive model. When range only Kriging estimates were calculated, they were done so to be able to model potential anisotropic random field. By considering bearing one can potentially leverage a greater number of measurements when the radial variogram shows low variability over increasing arc length distance. Just as with the range kriged estimate a radial stochastic estimate is produced. These can be combined by scaling the mean estimates by the inverse of the opposing variances.

There are a couple of benefits to this approach. First, when close to the center reference point, one might expect that signal statistics might be isotropic. This can be confirmed by the radial variogram which would have a minimal constant value. In this case, calculation of the mean can include a greater number of measurements (potentially from around the entire circle) that wouldn’t normally have been used in a purely range-based approach. Second,

anisotropic behavior can be identified by the increasing slope of the radial variogram. When this is detected, points that are in that lag-distance from the POI can be removed from consideration in calculating the mean and variance.

6.2 LCM Definition

The Local Connectivity Map (LCM) is defined as random field within a bounded space (Ω) in \mathbb{R}^d . It is referenced to a local reference frame. For example, for a mobile agent, the random field moves with the agent, and assuming a fixed mounted transceiver, the motion of the random field is tied to the pose and orientation of the mobile agent. Of course, for a fixed node (for example the static communications buoy previously described, there exists a unique pose and orientation (x_c) that describes the fixed communication node's position. Relative to this local reference frame, we seek mean and variance signal statistic estimates for a discrete number of points of interest within (Ω) . A local grid lattice is superimposed over the region where the dimension of each lattice cell corresponds to the map resolution where each point $(x_{(1,1)}^c \cdots x_{(n,m)}^c)$ is the centroid of the region. Figure (6.6) gives an example in the two-dimensional case.

In order to build the LCM, there are five necessary steps: data transformation, variogram analysis, solving a linear system of equations for λ , calculating the mean prediction, and calculating the variance. Strengths of the technique are the calculation of map variance and the straightforward solving of a system of linear equations for mean and variance estimates. Weaknesses of the approach include dependence on summarizing the empirical semi-variogram as a smooth function and the computational burdens of taking the matrix inverse for the linear system of equations.

6.3 Considerations for Communications

Wireless communications can be categorized in terms of narrowband and wideband systems. This is determined by the ratio of the bandwidth with respect to the center frequency. For example, IEEE 802.11 wireless systems are considered to be narrowband systems and underwater acoustic systems are generally considered wideband. There are several measures that are used to characterize the communications channel. First and foremost, received Signal to Noise Ratio (SNR) is an important parameter for determining the quality of the channel. Let P_T be the power of the transmission. Let P_r be the power of the measured at the receiver and P_w be the power of the noise measured at approximate the time

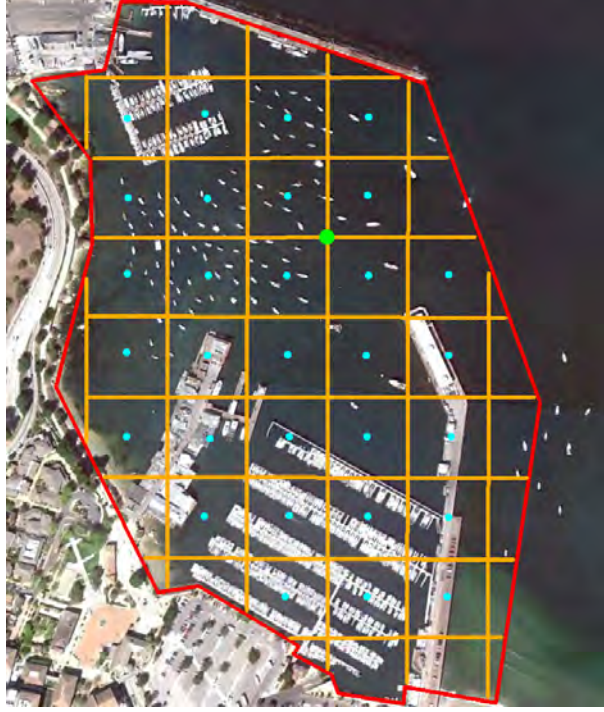


Figure 6.6: *LCM Example for Monterey Harbor, CA: The green dot represents a fixed reference point and is the source for all communication x_c . The boundary of the survey area (Ω) is drawn in red. The overlaid orange lattice represents the map resolution. This has been drawn disproportionately large for the graphic. The light blue dot is the centroid of the lattice cell where the mean estimate and variance is calculated but is assumed to hold inside the entire cell.*

of the received transmission. SNR is the ratio $\frac{P_r}{P_w}$. SNR is frequently displayed using the logarithmic decibel (dB) scale.

$$SNR_{dB} = 10 \log_{10} \left(\frac{P_r}{P_w} \right) \quad (6.24)$$

Received Signal Strength (RSS) is the principle measurement used in the dissertation for creating the connectivity maps. Unlike SNR, It doesn't consider the noise in the environment and is a useful metric since it doesn't have to consider the impact of a "moving floor". This is the potentially random way in which the SNR may change due to non-stationary environmental noise.

It is recognized that these are not the only tools for assessment of the channel quality.

Other metrics include the Bit Error Rates, Packet Error Rates and others could be used to provide a potentially better solution. The simplicity and availability of RSS makes it the most convenient metric to use.

Characterization of the communication channel is typically difficult due to the blocking, scattering, reflection and diffraction of the signal wave. The communications channel can be modeled as a function where there are three effects associated with the channel quality. They are small-scale or multi-path fading, large-scale fading and path-loss.

Small scale or multi-path fading results from multiple copies of the signal arriving at the receiver and are a result of the combined effects of signal attenuation, delay and Doppler shifts. Large-scale fading is due to signal collisions with objects and path-loss is the attenuation of the signal over distance. It is a conversion of a position of the signal power to heat. Overall, the combination of the effects can be represented with the following equation [132]:

$$\underbrace{\ln P(\theta, t)_r}_{\text{Received power}} = \underbrace{\ln P_T}_{\text{Transmitted power}} + \underbrace{\beta - \alpha \ln d(\theta, t)}_{\text{Path loss} \leq 0} + \underbrace{\sum_i r_i(\theta, t) n_i(\theta, t)}_{\text{Shadowing effect} \leq 0} + \underbrace{w(\theta, t)}_{\text{Multi-path fading and noise}} \quad (6.25)$$

where the pair (θ, t) is a transmission ray θ at time t , α and β are constants, $d(\theta, t)$ is the distance between the transmitter and receiver across a particular transmission ray, r_i is the distance traveled along the i^{th} object along the (θ, t) ray and n_i is the decay rate of the signal within the i^{th} object [48].

Figure 6.7 shows an example plot of the Received Signal Strength over increasing distance.

6.4 Undersea Acoustic Communications

There are several properties that distinguish undersea acoustic communications from wireless radio communications. They relate to the nature of the acoustic signal and the ocean's environmental properties. They include the following:

1. Acoustic signal attenuation or path loss strongly increases with signal frequency. This is due to absorption and spreading loss. The result is that, for longer distances,



Figure 6.7: An example of decreasing WiFi Received Signal Strength as a distance increases from the Mostofi paper entitled “Compressed Cooperative Sensing and Mapping in Mobile Networks,” [1].

- bandwidth is severely limited (approximately 1kHz [104]). For shorter distances the bandwidth improves but is limited by the size of the transducer. As a result of the limitations associated with undersea acoustic communications, multi-hop networks are more energy efficient and enable higher overall bit rates. This has direct implications for the use of AUVs as an undersea mobile communications network.
2. The speed of sound underwater is approximately 1500 meters per second but varies according to oceanographic considerations including water temperature, bottom type, surface wave conditions, wind, salinity and others. This has an implication for mobile undersea nodes attempting to communicate. The impact of a Doppler shift is proportional to the relative velocity ratio between two vehicles and the speed of sound. For mobile nodes communicating with ground-based radio systems, this ratio is small and the resultant Doppler shift can be ignored. This is not the case for mobile underwater nodes. The Doppler shift can create signal distortion between vehicles when the relative velocity is as small as several meters per second.
 3. There is a strong multi-path component to undersea acoustic communication. Sound waves are refracted from the ocean surface, floor, and objects. The sound propagation can also be channelized from temperature isoclines. Sound waves obey Snell’s Law and move towards regions of lower propagation speed. A particular path of an

acoustic ray may be longer yet arrive at the receiver more quickly. The implication is that techniques at the receiver must account for this statistical variability.

4. Noise characteristics are non-stationary and non-Gaussian. This is especially true in restricted harbor areas. Major components of noise include surface waves (strongly correlated with surface winds), rain, biological sources (such as snapping shrimp) and man-made (shipping).

6.5 Conclusion

This chapter focused on the development of the Kriging equations with an emphasis on the use of these equations in communication channel estimation. After deriving the ordinary Kriging equations an example was provided describing the steps necessary to create a mean and variance estimate of a quasi-stationary random field based on RSS simulated measurements. The procedure focused on a kind of additive Kriging model that used both a range and bearing variogram. This is a unique approach that is particular to communication and sensing devices that may propagate signals omni-directionally. The chapter defined the Local Connectivity Map as a random field and described considerations for communications and in particular undersea acoustic communications.

CHAPTER 7:

Experimental Results and Analysis

In the last chapter, a Local Connectivity Map (LCM) was defined and, through simulated data, a stochastic additive Kriging procedure was demonstrated. This chapter describes the results from the collection of undersea acoustic modem data. While the methodology in the dissertation has been developed for a general wireless communication, the application has focused on the undersea environment since it is believed to be, possibly, the toughest example of multi-node, multi-hop wireless networking. After presenting examples of the LCM, there is an analytical section for assessing the model's accuracy through cross-validation.

7.1 Data Collection

Signal measurements were collected in Monterey Harbor, Monterey, CA in August 2011 and July 2012. Figure 7.1 shows the five locations (S_1, S_2, S_3, S_4, S_5) of the Woods Hole Oceanographic Institution (WHOI) gateway buoy. It had an WHOI undersea acoustic modem hanging approximately four meters below the ocean surface. Additionally, there was a free-wave radio for monitoring the equipment and relaying data. A small Rigid Hull Inflatable Boat (RHIB) powered by a 250 HP water jet engine maneuvered around the bay and received the acoustic transmissions. The modem deployed from the boat was also submerged in approximately four meters of water.

The WHOI Micro-Modem [111] was configured at a nominal 27kHz, using Quadrature Phase Shift Keying (PSK) (as opposed to Frequent Shift Keying (FSK)) with a set baud rate of 5.3 Kbps with a 5000Hz bandwidth. The two modems provide half-duplex communication; meaning that the modems cannot broadcast and receive transmissions simultaneously. The gateway buoy was the source of the modem transmissions. The modem deployed from the boat was the destination of the transmissions. For these series of tests, a standard mini-packet was transmitted [133]. Useful information in the mini-packet includes the following:

1. SNR In: Input SNR, dB, channel one only
2. SNR Out: SNR at the output of the equalizer, dB
3. Symbol SNR: SNR measurement after despreading



Figure 7.1: The annotated positions in the image (S_1, S_2, S_3, S_4, S_5) are the locations of the gateway buoy for the first two rounds of testing (August 2011 and July 2012). For each dataset, the buoy (shown on the right-hand side) was placed at one of the locations and signal statistics were collected using this buoy as the source of a message packet where the destination was from a boat maneuvering around the harbor.

4. RSS: Received Signal Strength in dB (re. $1 \mu\text{Pa}$ at one meter)
5. Noise: Noise level measurement, taken before the start of the receipt of the packet, channel one only
6. MFD: Matched Filter Detector power level (another indication of signal strength indicator)

Figure 7.2 depicts the difference between the Input, Output and Symbol SNR measurements. The input SNR measures the SNR prior to the equalizer. It may be as low as 0 dB and the packet may still decode correctly. Maximum SNR is 30 dB. The highest Input SNR in the harbor data sets was 28 dB. The output SNR is measured after the signal has passed through the equalizer and therefore compensates for the multi-path. Greater levels of multi-path degrades performance. Symbol SNR is measured after de-spreading.

It is possible to build the LCM from SNR In, SNR Out, Symbol SNR and RSS. The difficulty with SNR is that the ratio is impacted by the unpredictability of changing noise

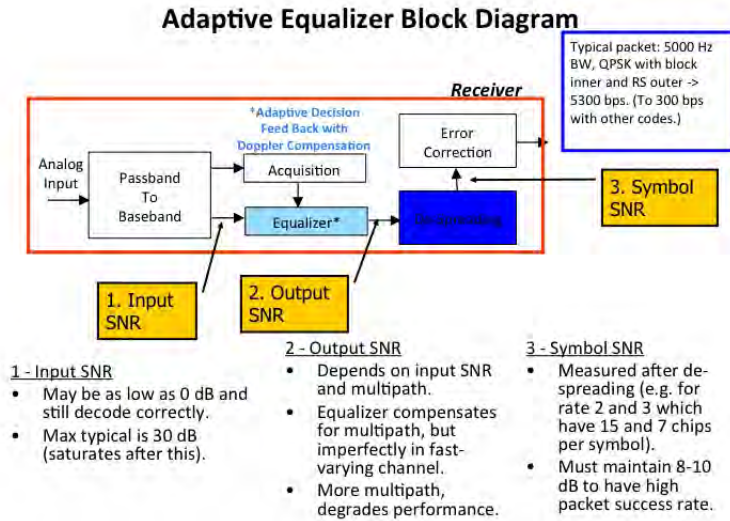


Figure 7.2: Description of the differences between the Input SNR, Output SNR and Symbol SNR with the WHOI Micro-modem. Graphic courtesy of Lee Freitag WHOI Principal Scientist.

Destination	Number of SNR measurements	Buoy Lat/Lon	Start time (PST)
S_1	1113	36.607024, -121.889747	1919
S_2	1298	36.606924, -121.891878	1810
S_3	1275	36.607755, -121.893768	1642
S_4	948	36.605322, -121.893274	1536
S_5	1274	36.604689, -121.889968	1449

Table 7.1: SNR measurements to each destination (August 2011)

levels. This is analogous to estimating altitude from a moving floor. For this reason and to be consistent with current research in (Radio Frequency Identification) RFID, indoor, (Global System for Mobile Communications) GSM and wireless localization, RSS is used for building the LCM.

7.1.1 August 2011

In August 2011, the surface boat drove patterns away from and toward the gateway buoy at a several different bearings. The objective was to continue driving away from the buoy until acoustic communications was no longer received or a boundary was reached. Table 7.1 shows the number of measurements relative to a single buoy. The data set is useful for giving a sense of the variability of the harbor measurements as a function of increasing



Figure 7.3: *The RSS measurements relative to the S_5 buoy position.*

distance. What follows is a brief discuss of the 5 datasets and some points to consider regarding the nature of the acoustic signals. The datasets are discussed in temporal order, that is, the S_5 was the first collected followed by S_4 , S_3 , S_2 and S_1 . All the harbor images are oriented so that a North direction is up.

Figure 7.3, shows the RSS measurements relative to the S_5 location of the static buoy. There are a couple of points of interest. First, even though the water is fairly deep (10 meters) it is a relatively poor area for communications. Typically there is an area in close proximity to the static buoy where RSS values are at their strongest. With the exception of a couple of measurements, the RSS is very weak close to the buoy and this was fairly atypical. In fact the signal gets better outside of an approximately 75 meter radius from the buoy position. There are many possible explanations for the reasons (muddy bottom of the ocean floor, transducer orientation), but this poor signal attenuation was consistent with all measurements (over multiple days) from this region.

Figure 7.4a, shows the RSS measurements relative to the static buoy located at S_4 . It shows the ability of the signal to penetrate through a dense pylon “forest” that supports commercial pier. There is also quite a bit of variation in close proximity to the static buoy. This highlights the difficulty in using automated Kriging techniques for creating an estimate

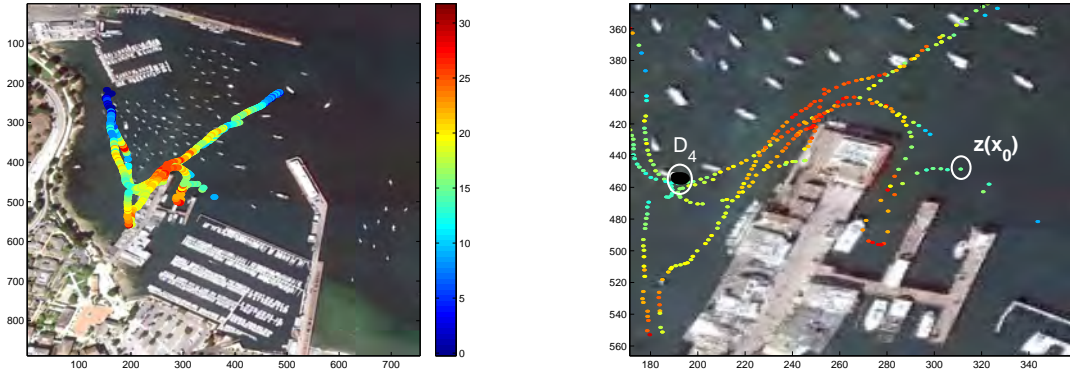


Figure 7.4: The RSS measurements relative to the S_4 buoy position. On the right-hand side show a magnification around the commercial wharf.

of signal strength. Figure 7.4b shows a closeup of the area. The way in which points are selected either for the semi-variogram of the estimation process for a point of interest could dramatically change the Kriging estimate.

It also highlights a potential strength of communication-centric Kriging technique—given *a priori* measurements, one can deliberately select a bearing window to construct a semi-variogram model that is appropriate given the data. In other words, it can handle the anisotropic nature of the data regardless of obstacles that inhibit communications. Acoustic communications between the fixed node S_4 and location x (that produced the measurement $z(x)$) are degraded but still possible through the pier and this would be difficult to predict this ahead of time.

Notice that the boat took two bearings away from the buoy. One bearing hugged the coast and the other went out into the center of the harbor. The data shows that the degradation of signal strength is anisotropic—it is a function of bearing. In this case, a possible hidden explanation is the shallow water along the shore attenuated the signal. Finally, the measurements taken along the bearing out into the center of the harbor show an increase in signal strength approximately 50–75 meters from the buoy followed at a degradation over range.

In Figure 7.5, the buoy was located in shallow water (3 meters) and the signal attenuated quickly such that the ability to communicate (assuming a minimum RSS level of 8–10dB to successfully pass acoustic messages) was limited to approximately 300 meters. What is

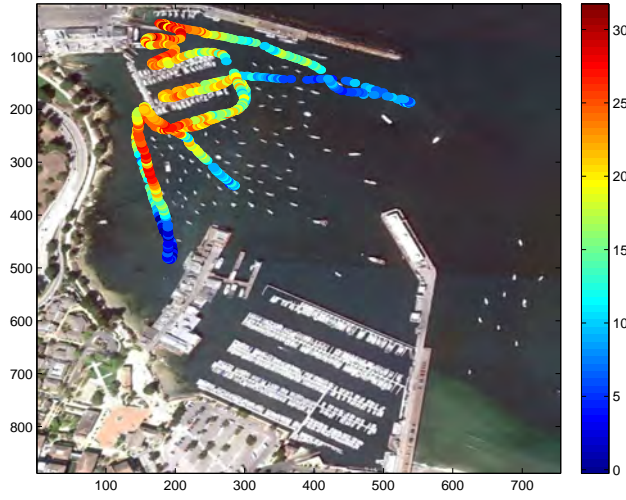


Figure 7.5: *The RSS measurements relative to the S_3 buoy position.*

interesting about the dataset was the strong signal that penetrated the natural and manmade objects located in the upper left-hand region. Given a copy of the map one could not assume that it would be possible to communicate in this region, since one could not be sure regarding the underwater infrastructure associated with the docking in the area. This is an example of the potential utility of a data-driven approach where a system of vehicles can exploit the ability to communicate that permits a less conservative collaborative approach than would otherwise be warranted.

In Figure 7.6, the buoy is close to the center of the harbor. There are many surrounding moored boats in this region and the RSS measurements are very dependent on the bearing, but overall, the dataset shows lower levels of signal attenuation in nearly all directions when compared with the S_5 , S_4 and S_3 . A couple of points worth noting. Near the commercial pier, along the southwest bearing, the RSS signal is somewhat consistently different between going away versus going towards the buoy. Although this was seen in several different data sets it was somewhat surprising that it was not detected more frequently. Second, with the measurements heading north by northwest, at the northern most measurements there was a significant improvement in RSS and this unexpected area of stronger signal strength reflects the challenges associated with purely physics modeling for channel estimation. There were a total of 15 measurements in a 15 meter radius where the mean

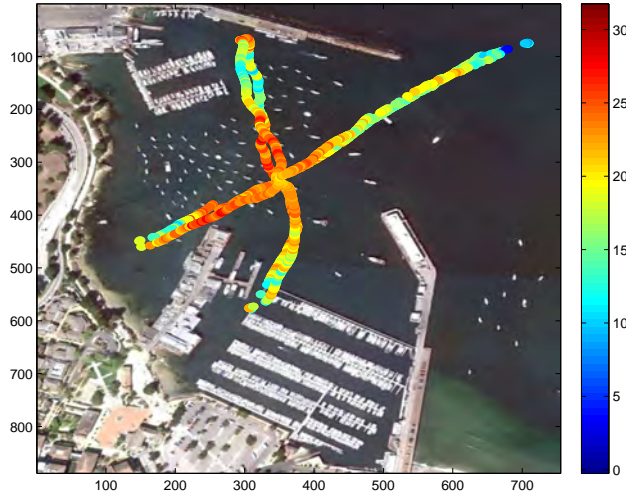


Figure 7.6: *The RSS measurements relative to the S_2 buoy position.*

RSS value was 24dB (The maximum value of RSS seen in the harbor was 28dB).

Like the previous dataset, S_1 has examples where the RSS measurements increased at the farthest points from the buoy both in the north and east directions. Otherwise, the dataset is characterized by a uniformity of measurements except to the south around the municipal pier. This is maybe attributable to the water depth at the buoy location (17 meters).

In a sense, this is the “best” data set from the perspective of being the least impacted by obstacles. Still it is possible to look at the differences in signal strength over range. Figure 7.8 show the S_1 data set partitioned into four subsets annotated with red, green, blue and magenta triangles. Figures 7.9a-d shows the corresponding range versus received signal strength plots in the matching color. This highlights the isotropic (or anisotropic behavior) by comparing different bearing plots. Each plot has a best fitting fourth-degree polynomial included with the data. It is interesting to note that there is a similarity in the northeast and southwest plots. As the range increases the data becomes more variable. For the northwest and southeast plots the data remains more concentrated as the range increases.

Additionally, the next day, an AUV was deployed in the harbor at the five buoy positions. It remained roughly in a single position using cross-body vertical and horizontal thrusters. Table 7.2 gives a summary of the oceanographic data at each of the sites.

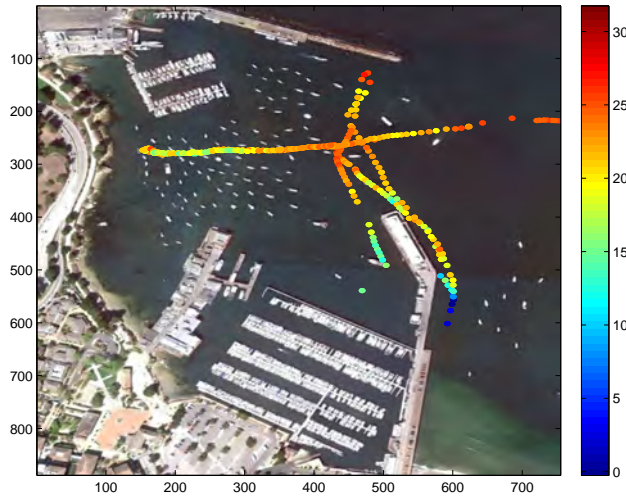


Figure 7.7: The RSS measurements relative to the S_1 buoy position.

	S_1	S_2	S_3	S_4	S_5
Time	22:05:44.6	21:55:20.0	21:44:54.0	21:29:27.9	21:04:49.2
AUV Depth (m)	7.5	6.0	1.0	1.0	4.5
Bathymetry (m)	10.1	8.35	2.1	2.3	6.75
Water Vel (cm/sec)	15.0	15.0	20.0	20.0	12.5
Water Directions (deg)	75/255	75/255	75/255	75/255	75/255
Salinity (gms/liter)	31.75	31.5	31.55	31.56	30.625
Sound Speed (m/sec)	1497	1497	1502	1501.8	1498.5
Temperature (c)	13.25	13.25	15.03	14.74	14.23
Latitude	36N36.421	36N36.414	36N36.459	36N36.315	36N36.286
Longitude	121W53.433	121W53.519	121W53.630	121W53.597	121W53.398

Table 7.2: Oceanographic data collected from an AUV in Monterey Harbor, CA (August 2011)

7.1.2 July 2012

The dataset in 2011 was useful in getting a better understanding for how the acoustic signal propagates in a relatively confined environment. When compared with a more open ocean environment, there was a stronger attenuation of the signal in the harbor. This knowledge is useful for determining appropriate spacings for maintaining a network of mobile collaborating AUVs. A limitation of the data collection was the sparse coverage. The measurements were confined largely to straight out and back bearing lines relative to the static buoy.

In July 2012, another data set was collected. The goal this time was to more completely cover the harbor region relative to the fixed gateway buoy. This was accomplished, when

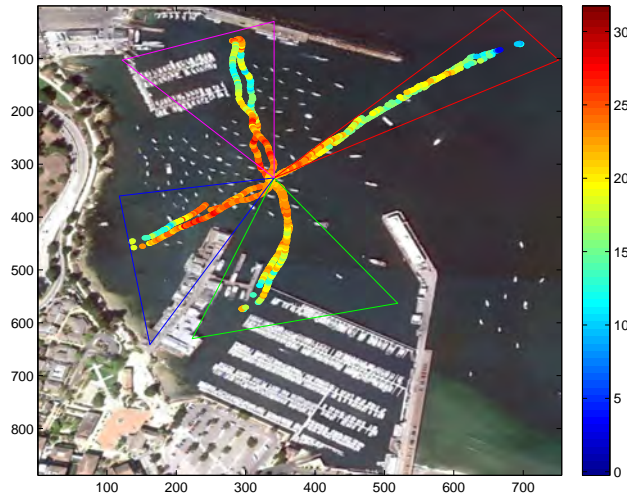


Figure 7.8: RSS measurements (relative to S_2 buoy) partitioned into 4 groupings annotated by the red, green blue and magenta triangles

Destination	Number of SNR measurements
S_1	963
S_2	752
S_3	525
S_4	500
S_5	589

Table 7.3: RSS Measurements to each destination (July 2012)

possible, using a mow-the-lawn pattern. With the greater coverage, it provided an opportunity to construct LCMs with lower overall variance since there were fewer spatial gaps in measurements. Table 7.3 shows the number of measurements for each of the buoy locations and the start time for each dataset.

Again the data is presented in temporal order and was collected over the course of a single day. Figure 7.10 gives the tides for the day. Figure 7.11 shows the RSS measurements relative to the buoy located at S_1 . The additional measurements provide a clearer mapping of signal strength. In particular the center and entrance to the harbor show a strong signal. It drops off substantially along the shallow water area along the western shoreline. There is mixed signal strength in the southern region where, at one point, it dropped off quite a

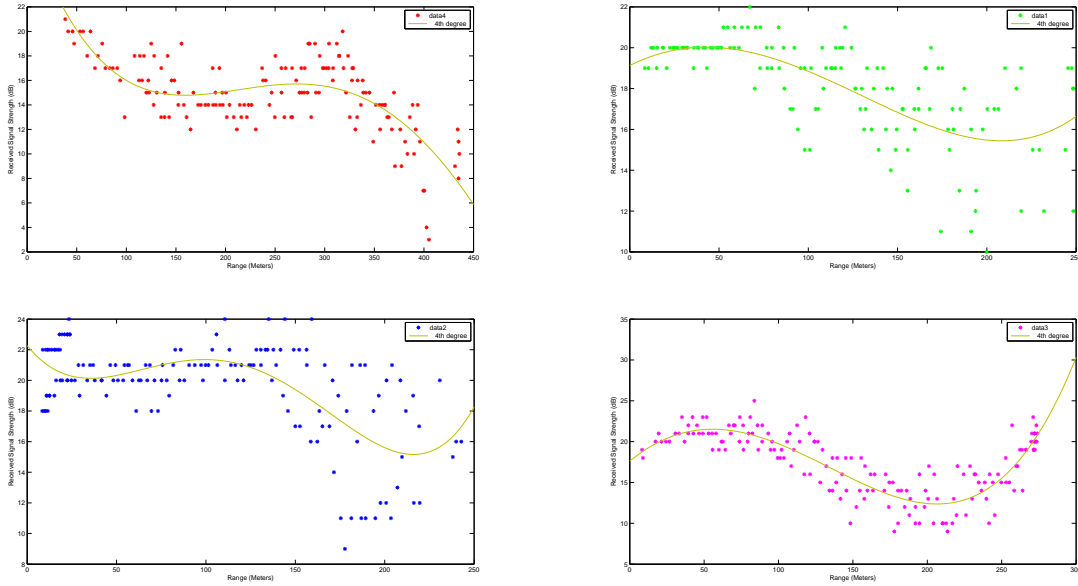


Figure 7.9: RSS plots and best fitting polynomial functions corresponding to the red, green, blue and magenta triangles from Figure 7.8

bit (at local position [500,500]) and as the boat returned a second time the signal was much stronger.

Figure 7.12 provides the RSS signal strength relative to the buoy located at S_2 . In general it shows a surprising consistent signal over almost the entire area. The exception is the southern area (local position [500,500]) where the signal drops off substantially. Also of note is the northern region where the signal began to drop off (350,150), but increases again near the Coast Guard pier. This is unusual since the image doesn't show a newly built dock with pylons and boats (at approximate position [300,100]) that one might expect would attenuate the signal.

The S_3 buoy position (Figure 7.13) is tucked into a shallow water area that is occluded by boat docks. The acoustic modem transducer was located approximately 2 meters underwater. The RSS measurements relative to this modem show that there was a reasonable signal in close proximity around the docks but the signal was degraded significantly in the southern section of the harbor. What is particularly interesting was a band of strong signal in the eastern direction—it reflects a possible ducting or channeling of the signal that was not greatly impacted by the docks.

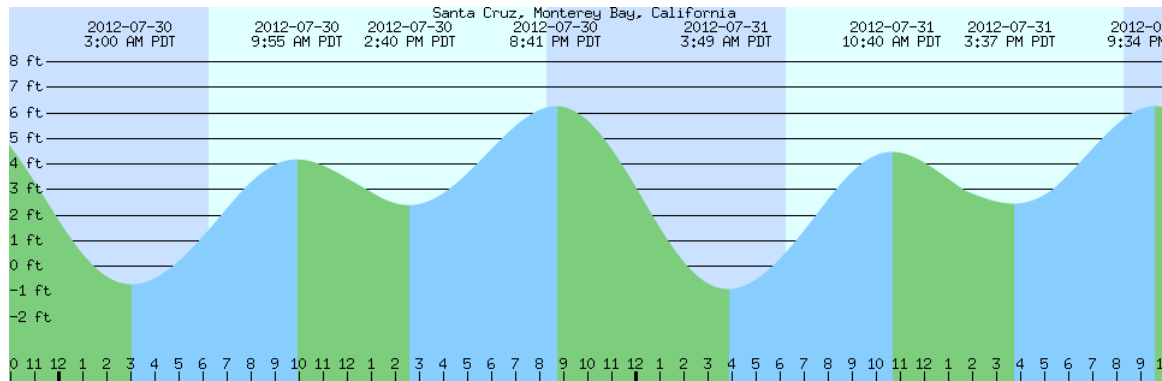


Figure 7.10: *Tides for July 30, 2012*

Figure 7.14 shows the plot of the RSS measurements relative to the static buoy at the S_4 position. The buoy is also tucked into a confined area in which the signals could be occluded by the commercial wharf. The water is a bit deeper here and combined together with a more direct path to the central part of the harbor may have contributed to a stronger signal over a larger area of the harbor than when compared to S_3 .

When compared to the other datasets, the RSS measurements relative to the S_5 buoy given in Figure 7.15 are maybe the most interesting. As can be seen in the plot the signal is almost uniformly poor no matter what the range from the buoy. This is surprising since the water is relative deep and the area is open such that there is a direct path to much of the rest of the harbor.

In summary, the data sets provide insight into undersea acoustics in a confined shallow water, harbor environment. It is one of the first datasets collected for investigating the acoustic attenuation in confined areas. It is also somewhat unique in that data was collected by mobile platform. Normally undersea acoustic measurements are conducted from pairs of static nodes. The next section takes these datasets and applies the Kriging techniques that were developed in the previous chapters. By using the Communication Kriging techniques at all grid points within harbor region and within 300 meters of the source node, the result is an LCM.



Figure 7.11: RSS measurements relative to the S_1 buoy position. July 2012

7.2 LCM Results

The section describes the results of using the July 2012 datasets to build LCMs using ordinary Kriging interpolation. The map resolution was set to 10 meters—meaning that Kriging was used for estimating a point of interest located in the center of the grid cell and it is assumed that the value is constant within the cell boundary. There are two components for each gridded cell the RSS mean estimate and the variance. Figures 7.16-7.20 present the results.

7.3 LCM Accuracy

LCM accuracy is impacted through three main sources of errors: Sensor-related, location estimation and modeling. Sensor measurement errors are associated with the ability of the sensor to accurately measure the physical environment. In this case, sensor measurement error reflects the ability of the transceiver to accurately assess the acoustic signal. With the WHOI μ -modem it reflects the ability of the transceiver to measure the ambient noise in the environment and the received signal strength (RSS). The RSS is considered the more reliable of the two measurements since it is measured over the entire length of the incoming



Figure 7.12: *RSS measurements relative to the S_2 buoy position. July 2012*

packet. Conversely, the noise measurement is made over a small window due to the modem constraints of real-time operations. ⁶

The noise measurements from the five datasets collected in July 2012 varied from -70 to -48 dB. As to be expected, these noise levels are considered very high in comparison with open ocean water. Recent testing in a Scotland Loch produced nominal noise levels of 30dB. ⁷ The high noise levels in Monterey Harbor (50–60 dB) were not surprising since the measurements were taken from a noisy boat in an active harbor.

Figure 7.21 shows an example of the raw noise measurements. It shows a dispersed pattern of high noise levels that is interspersed with fewer low noise measurements. A nice aspect of the noise data is that, unlike the signal data, it is not referenced to a central source point; the measurements were taken by the acoustic modem as the boat moved through the harbor. Because of this, the sequence of noise datasets taken throughout the day may give insight into the spatial-temporal process.

A final note of noise estimation. It too could be a useful measurement for determining

⁶Written correspondence with Lee Freitag (December 2012), WHOI Senior Researcher and developer of the μ -modem.

⁷Written correspondence with Lee Freitag (December 2012).



Figure 7.13: *RSS measurements relative to the S_3 buoy position. July 2012*

when to communicate. The ability to build a map for determining either regions or times of decreased noise should improve the probability of successful communications between vehicles. While not a focus for this work, the combination of separate signal and noise maps used conjunctively could potentially be a useful future technique.

The ability to estimate RSS at a particular location is, of course, impacted by the ability to determine the location at which the measurement is taken. In many instances this is not a problem since GPS is widely available. This is not the case for underwater, indoor and space. The ability to accurately estimate position adds an extra level of complexity to estimating the spatial random field. In essence, any uncertainty associated with the position adds to the uncertainty of the random field.

7.4 Model Accuracy

In assessing the quality of the LCM, it is necessary to consider issues of prediction accuracy. This is the ability of the Kriging techniques to accurately interpolate data to provide minimal errors at locations where there are no direct measurements. The technique used provide an assessment of the accuracy of the Kriging techniques is known as K-folds cross-validation.



Figure 7.14: *RSS measurements relative to the S_4 buoy position. July 2012*

7.4.1 Cross-Validation

The basic idea with cross-validation is to partition a data set such that a large percentage of the data is used to build the model and the remainder is used to check model validity. The validation process takes each of the withheld measurements z_{wh} and compares them to the nearest neighbor prediction point in the model $Z(x_{nn})$. The difference between the measurement and the prediction is the Kriging error (ϵ_{ke}):

$$\epsilon_{ke} = z_{wh} - Z(x_{nn}) \quad (7.1)$$

The standardized error scales the Kriging error by the square root of the Kriging variance estimate at the point ($Z(x_{nn})$) such that $\epsilon_{se} = \epsilon_{ke} / \sigma_k$.

There are several techniques for determining how to divide the data into the model set and validation set. Two common techniques are known as K-fold and “leave one out cross-validation” (LOO-CV). K-fold divides the data into k equal-sized partitions. For each iteration, one of the partitions is removed from the total dataset and the remainder is used to build the model. Afterward it is validated with the removed members. This happens k

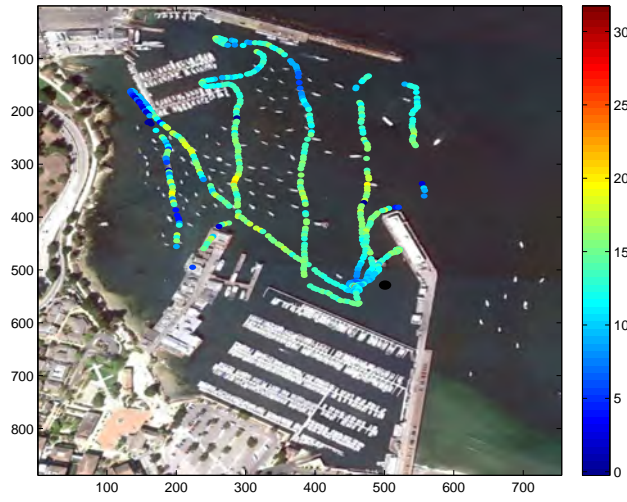


Figure 7.15: RSS measurements relative to the S_5 buoy position. July 2012

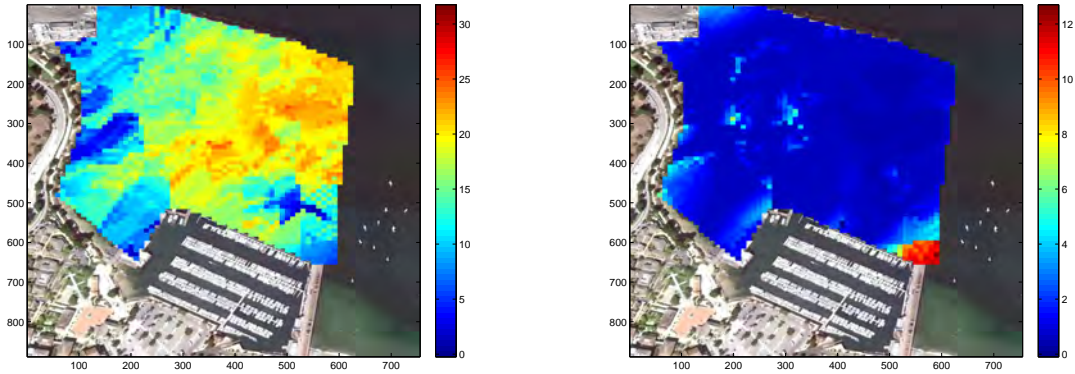


Figure 7.16: Kriging mean and variance results from the S_1 dataset

times and the results are summed in order to minimize any potential discrepancies due to a particular partition. LOO-CV takes the k-fold methodology to its extreme where the model is built with all but a single data point and the remaining data point is validated against the model and this is repeated for the number of elements in the dataset. This analysis uses the k-fold cross validation technique with ($k = 10$).

There are several plots that provide insight into the Kriging model. First, given a theoretical

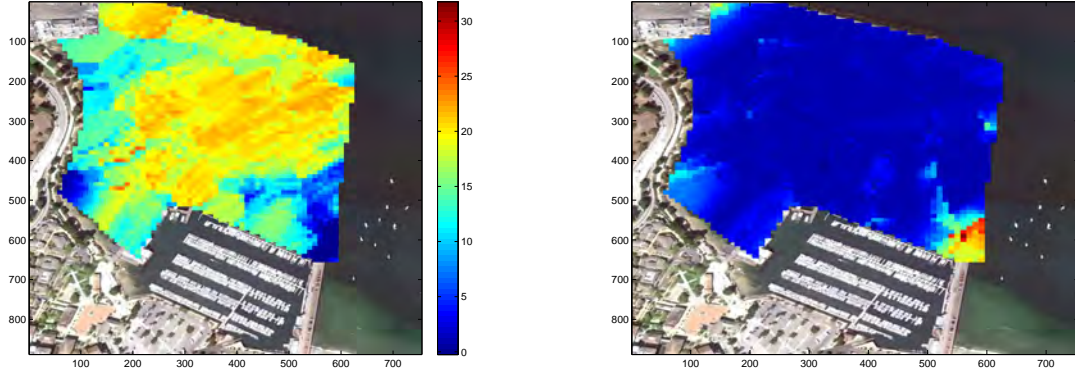


Figure 7.17: *Kriging mean and variance results from the S_2 dataset*

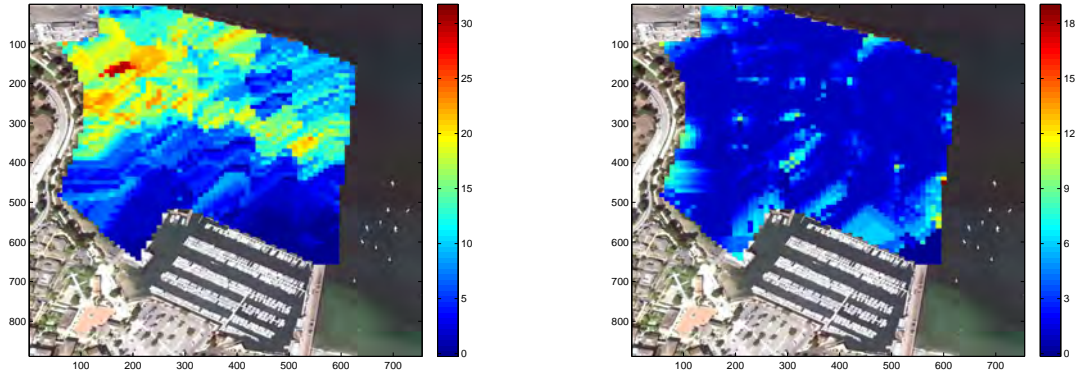


Figure 7.18: *Kriging mean and variance results from the S_3 dataset*

variogram, the Kriging error should have a zero mean $E[\varepsilon_{ke}] = 0$ Gaussian distribution with variance σ_K^2 . This can be graphically checked by looking at the Kriging error histogram of the aggregated k-fold results. Figure 7.22 shows histograms for each of the S_1 through S_5 data sets with a fitted Gaussian distribution fitted to the data. Table 7.4 gives the Kriging error mean and variance of the best fitting Gaussian distribution. As can be seen in the histograms, they generally conform to a Gaussian distribution with the exception of an overabundant number of errors close to zero.

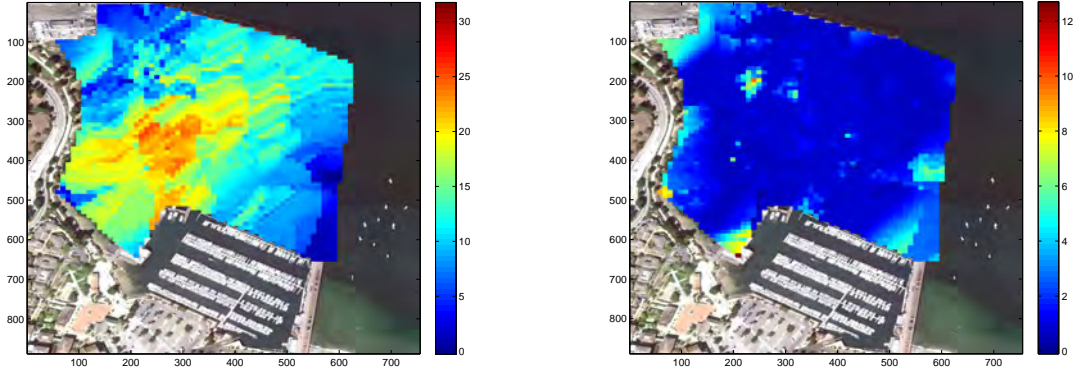


Figure 7.19: *Kriging mean and variance results from the S_4 dataset*

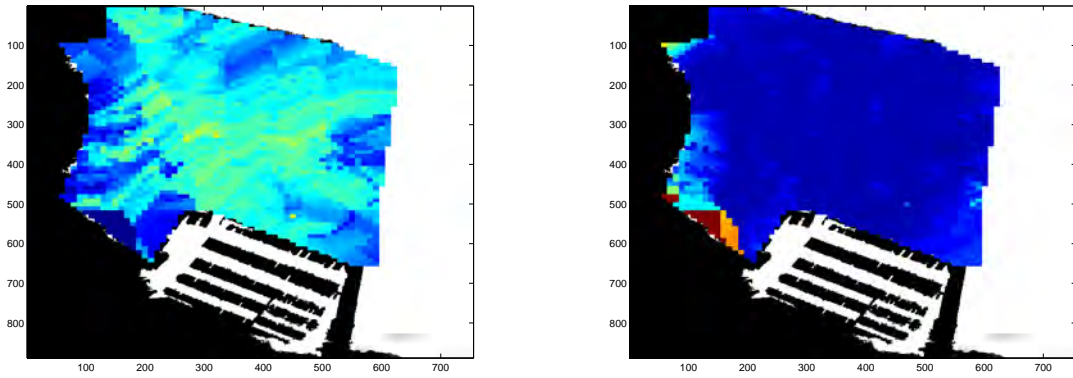


Figure 7.20: *Kriging mean and variance results from the S_5 dataset*

7.5 Conclusion

The chapter has presented a technique for creating a Local Connectivity Map (LCM) that can be used for estimating the RSS relative to a fixed transceiver. It is constructed using Kriging techniques for estimating the mean and variance of a Gaussian random field. Results were described based on underwater acoustic data collected Monterey Harbor, Monterey, CA. The results helped to illustrate the advantages of the technique and the anisotropic nature of communication channel estimation in a cluttered and dynamic area.

With the LCM in hand it is possible to use this information to more optimally position nodes in a configuration that would support greater network reliability by decreasing trans-

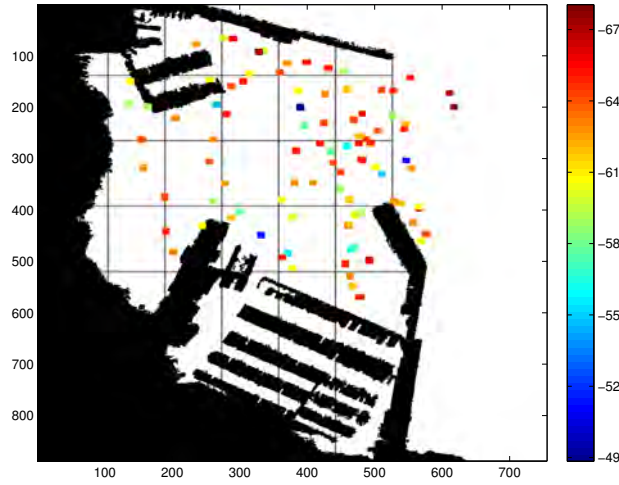


Figure 7.21: The image overlays the background noise (dB) on top of a binary representation of the Monterey Harbor. The color scale is the Matlab “jet” colormap.

Destination	Kriging error mean	Kriging error variance
S_1	-.002	5.24
S_2	-.003	4.42
S_4	-.059	5.50
S_5	-.094	6.13

Table 7.4: The table provides the mean and variance associated with a best fitting Gaussian distribution for the Kriging error (ϵ for the S_1 through S_5 data sets).

mission failures. There are several limitations that need to be addressed for a collaborative navigation implementation. First, the Kriging technique assumes complete *a priori* RSS measurements. In many circumstances this may not be available. What would be preferred is to build the LCM as data becomes available.

Second, the LCM provides an estimate from local reference frame. For distributed collaborative operations this is not sufficient. What is need is a methodology that is not bound to the pose of the mobile agent. In other words, what is required is the ability to produce an estimate of signal strength between any two points in a region so the collaborating agents can use this information in planning. This is addressed in the next section.

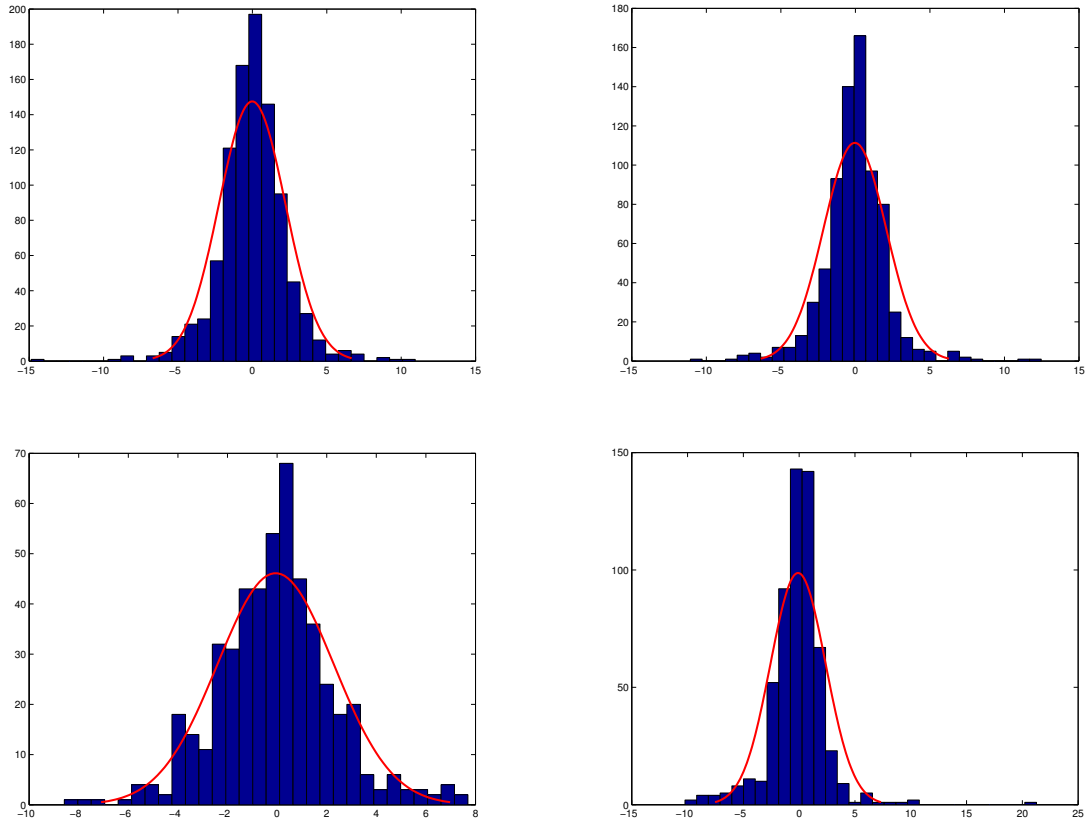


Figure 7.22: The aggregated Kriging error for the datasets (S_1, S_2, S_4, S_5) using the k -fold cross validation where $k=10$. The x-axis is the Kriging errors divided into bins of equal size and the y-axis is the number of elements in each bin.

CHAPTER 8:

The Global Connectivity Map

Part one focused on the LCM. It predicted mean and variance RSS at designated points of interest relative to an agent's local reference frame. While LCMs are useful in determining positions of a static network of communication nodes or maintaining a formation with mobile agents, it is insufficient for a collaborative mobile team conducting a survey in a bounded area.

This is due to the fact that in a survey objective, there is a greater emphasis on the ability of a geographical region (versus an agents ability) to support wireless communications. To support collaborative mobile surveying, a candidate framework should have the following attributes:

- Assess the communications channel between two (or more) points located anywhere in the bounded region of interest.
- Incrementally build the model as measurements become available.
- Calculate mean and variance estimates.
- Build the model in a distributed, collaborative way.
- Build the model in a computationally efficient manner.
- Pass minimal information between agents to separately maintain up-to-date models.

This chapter introduces the Global Connectivity Map (GCM). Its purpose is to enable mobile agents to assess the quality of communication between any two points within a bounded region. It expands upon previous chapters, in that, the technique to build the GCM involves the simultaneous estimation of multiple LCMs. The use of the term "Global" reflects a convention in robotics signifying the connectivity map is anchored to a common reference frame (e.g., GPS) accessible by multiple mobile systems.

The reason for the introduction of the GCM is several-fold. First it is a natural outgrowth of the LCM. In other words, it is consistent with the concepts of a local and global framework that is used to relate mapping and navigation techniques within the robotics literature. Second, and more importantly, the GCM can be a key technology for improving multi-vehicle navigation. Given a accurate GCM, it could improve collaborative navigation algorithms

by fundamentally providing a better model regarding communication assumptions between vehicles. This can improve system robustness, flexibility and effectiveness. Some examples include:

- An agent can temporarily move out of communication range with the knowledge of where to return for connecting back into the network.
- An area can be partitioned, taking into consideration the communication landscape, such that it optimizes the probability that networked agents remain fully connected.
- Development of trajectories (path planning with time constraints) between a system of vehicles such that area coverage constraints are obeyed.

The process to build the GCM, is a generalization of the LCM methodology. With the GCM, a spatial tessellation is defined such that there is a centroid position associated with each tile and from each center we seek to estimate an LCM (a random field).⁸ A full realization of the tessellation of random fields that completely covers the bounded area permits communication channel estimation since each vehicle is within the bounded area and knowledge of their respective position allows one vehicle to predict the ability to successfully send wireless communications to one or more other vehicles.

With the LCM, there was an assumption of quasi-stationarity. Any random variable, within a defined neighborhood, has the same mean value. With the GCM, the quasi-stationarity assumption holds for random *fields* such that within the boundary of a GCM tile, all random fields are homogeneous such that they can be represented by a single random field. For standardization, this single random field is located in the center of the tile.

At this point, it may be useful to review an envisioned concept of operations. A group of unmanned, wireless, mobile vehicles are tasked with conducting a detailed survey of an area. There is limited information available ahead of time, so any prior assumption must be carefully considered. The system of vehicles depend on wireless communications for coordinated behavior and cooperative navigation algorithms for bounding the position errors and conducting the survey.

The initial objective is a rapid reconnaissance of the area in order to assess the environment. There are several objectives within this rapid assessment and one of them is to assess the

⁸A tile is defined as a unit of the tessellation where it is not necessarily limited to a rectangular shape but can include irregular and dissimilar shapes (e.g. a Voronoi tessellation or honeycomb tile).

ability to communicate—to build the GCM.⁹ It is envisioned that it would be used to help partition the survey into sub-regions such that each vehicle is assigned a region for detailed surveying. The partitioning of the region is to ensure that, to the greatest degree possible, the group of robots are able to effectively communicate. This is the focus of this second part of the dissertation.

With the pre-survey complete, the main survey is conducted. Each vehicle is assigned a region for detailed survey. As the vehicles communicate, they use the GCM to determine when and where to transmit and receive. This communication is, of course, used to continually improve the GCM. This in turn may be used to help redefine boundaries as the survey continues. Future work will investigate the coordinated path planning problem which takes into account where and when to take measurements to realize the GCM.

There are many techniques that can be used to build the GCM from a limited dataset. Many rely on some form of Bayesian inferencing. The reason is that it addresses many of the desired attributes listed above. Bayesian inferencing supports incremental updates of the GCM as well as batch updates. Second, it requires little in the way of prior assumptions. Third, it provides the ability to calculate the mean and variance of the random fields. Fourth, many of the techniques can be implemented in a distributed fashion and in a computationally tractable way.

8.1 Relevant Research

There has been extensive research into modeling Radio Frequency (RF) received signal strength. The main application is the localization of mobile agents for internet context-dependent, content delivery. A common approach is known as a fingerprint technique [134], [135], [136] and [137]. It consists of two phases; an offline training phase and an online localization phase. In the offline training phase a radio map is built by evaluating the RSS measurements at fixed access points. In the second phase, the RSS measurements are used for estimating location. There are number of researchers that have also considered the temporal aspect of the RF localization problem including [137], [138] and [139]. Cheng et al. evaluate the feasibility of building a wide-area WiFi-based positioning system [140].

⁹Note that in an application where multi-vehicle SLAM is required, another important consideration is the location of readily detectable features. The combination of features and communications are then used together for determining the partition borders.

These concepts are conceptually similar to the GCM—the goal is to build situational understanding about the nature of the signal strength relative to position, but they differ in three important aspects from the GCM. First, the goal for the GCM is to provide a signal model for establishing better collaborative control. The design goal in the other techniques was not for a control input. Second, with the RFID techniques, radio maps are built referencing fixed access points. This is analogous to the LCM presented in part one. For the GCM, there is no fixed access point, the goal is to estimate the ability to communicate as if there was an access point at the centroid position. Third, the goal is to build the GCM with a near real-time emphasis. This effectively blends the first and second phases of the fingerprint technique.

Another related research community that is interested in signal strength estimation involves cooperative localization in wireless networks [52], [91]. With cooperative localization, the idea is to maintain an estimate of the location of a network of sensors so the data collected from the sensors remains relevant. Again the methodologies are similar in terms of a measurement phase and a location-update phase. In the measurement phase, it involves exploiting RSS for the difference in signal loss between sender and receiver to determine the connectivity between nodes and the use of time of arrival (TOA) and time difference of arrival (TDOA) such that a combination of the approaches can be used for providing a more refined distance estimate. In the location-update phase, this information is used as input into a localization algorithm.

It is believed that the term “connectivity map” comes from the 2005 Kamakaris and Nickerson paper [141]. Cottingham describes the impact connectivity maps could have on vehicular traffic and includes techniques for estimating coverage maps (including Kriging) [34]. The principle estimation problem is the ability to predict signal strength with fast moving mobile nodes. As vehicles become driveless, access to communication conditions become increasingly important for at least two reasons. First, the collection and transfer of data to a central repository can be used by the driverless systems to improve performance and safety. Second, as vehicles become driverless, passengers will likely use mobile wireless devices more heavily. Reference for connectivity maps include, [33], [35] and [142].

This, I believe, is the first introduction of the concept of a Global Connectivity Map. Given a copy of the map it is possible to assess the ability to communicate between any two points in the bounded region defined by the map. The technique used to build the GCM

involves the simultaneous estimation of multiple random fields. The key underlying tenet for building the map is that instead of relying on a signal model the goal is build a model solely based on measurements. The strength of the approach is that it can easily take into consideration the vagaries associated with signal propagation and the environmental impact on the signal. In other words, the strength of the approach is particularly relevant for difficult environments, such as urban areas, indoors and underwater. Areas that are also dynamic and obstacle laden.

8.2 Overview

Part two of the dissertation is organized as follows. This chapter defines the GCM. The next chapter (Building the GCM) discusses a technique for building the GCM in near real-time. It is called a Spatial Communications Kalman Filter (SCKF) and it combines a Gaussian Process Model (GPM) with a recursive Kalman filter. The following chapter(Refining the Map) is a methodology for comparing neighboring random fields in the interest of simplifying the GCM. Part two concludes with examples of GCMs including accuracy estimates using acoustic modem data in Monterey Harbor. It also includes a comparative analysis between GCMs that reflect different data sampling strategies.

An emphasis is on reducing the number of measurements to build and maintain the GCM. There are at least four distinct strategies. The first is to make simplifying assumptions that minimally impact the accuracy of the model. With Bayesian inferencing techniques, one can use priors to speed the convergence of the model. The downside is that if these assumptions are incorrect it may take a longer period to converge than if the model started with no assumptions (a uniform prior).

The second is the use of a recursive filter instead of a batch algorithm permits near real-time estimation. Normally, in a traditional, supervised, machine learning algorithm, a complete data set is necessary for developing the model. Instead the model estimate is developed while the agents are conducting the survey. Third, while not developed in this dissertation, a collaborative navigation strategy can be developed such that the path for each vehicle collectively accounts for coverage to intelligently sample the environment and reduce total measurements.

Fourth, in Chapter Ten, a methodology using Probabilistic Graphical Models will be combined together with the Spatial Communications Kalman Filter for comparing random

fields in the spatial tessellation. This is a framework for a simple, common sense strategy. Suppose there are two center points for two random fields that are separated by a short distance. One might hypothesize that the two fields are similar. This could be used for redefining the tessellation boundaries. Given that two regions share similar signal characteristics the number of measurements necessary for estimating the GCM would be reduced since it is then plausible to use a single measurement for estimating the two random fields. As will be seen, this turns out to be a graph structure learning problem.

8.3 GCM Definition

The GCM is defined as a bounded survey space Ω in \mathbb{R}^d that is partitioned into a spatial tessellation ($\Omega = \cup_{i,j=1}^{n,m} \Omega_{nm}$). Within the boundary of a single tile Ω_{nm} , there exists a quasi-stationarity assumption for random fields such that a single random field can be used as single representation for the tile. From the center of each tile cell $\Omega_{nm}(x^c)$, we seek to predict the mean and variance of the random field $Z_{nm}(x), \forall \mathbf{x} \in \mathbb{R}$ out to a maximum distance R_{max} . Each tile in the tessellation has

- A scalar estimate of the overall ability to communicate omni-directionally.
- A scalar estimate of the overall uncertainty in the region.
- A discrete number of vectors where each corresponds to a bearing (relative to the LCM center) and each element in the vector corresponds to a mean and variance estimate at a particular distance from the LCM center.

Regarding the first bullet. The calculation of integral $\int_V Z(x)dx$ over some domain V in \mathbb{R}^d of a stochastic process is a common methodology for summarizing information content of a random field [143]. For a stationary random field, there has also been research on studying the behavior of integral predictors [128, 129]. Normally the set V is a square for geostatistics applications. For communications it is assumed to be a polar grid.

In an ideal situation, the measuring of the environment would encompass regularly spaced measurements along ranges and bearings. After the estimation process it is possible to calculate the integral of the continuous random field. The integral formula would be:

$$\int_0^{2\pi} \int_0^{R_{max}} Z_{nm}(\psi, r) r dr d\psi \quad (8.1)$$

where R_{max} is a defined maximum range and $x = (x_1, x_2)$ is point in \mathbb{R}^2 within the region R_{max} such that $x_1 = r \cos(\psi)$ and $x_2 = r \sin(\psi)$.

For a cluttered region, this is impractical it will not be always possible to fully measure the environment. The approach taken is to estimate the integral discretely using a double Riemann sum.

$$RS_{nm}(Z_{nm}) = \sum_{i=1}^v \sum_{j=1}^u (Z_{nm}(\psi_i, r_j^*)(r_j - r_{j-1})),$$

$$r_{j-1} \leq r_j$$

where $Z_{nm}(\psi_i, r_j^*)$ is an estimate of the random field mean at bearing ψ_i and range r_j^* . The estimate of uncertainty (the second bullet) is also calculated in similar fashion except that it is normalized with respect to the other cells. This ensures that the area of the cell does not influence the measure. Otherwise the scalar measure would be biased towards regions of lesser size. This is a measure of an uncertainty per unit squared where the unit is equivalent to the map resolution. It is given by the following equation:

$$RS_{nm}(\sigma_{nm}) = (\sum_{i=1}^v \sum_{j=1}^u (\sigma_{nm}(\psi_i, r_j^*)(r_j - r_{j-1}))/v),$$

$$r_{j-1} \leq r_j \quad (8.2)$$

$$v = \sum_{i=1}^n \sum_{j=1}^m RS_{ij}(\sigma_{ij})$$

An example of a one-dimensional Riemann sum is given in Figure 8.1, where the summation of the rectangles serve as an approximation of a Matérn function representing an approximation of the received signal strength as a function of range.

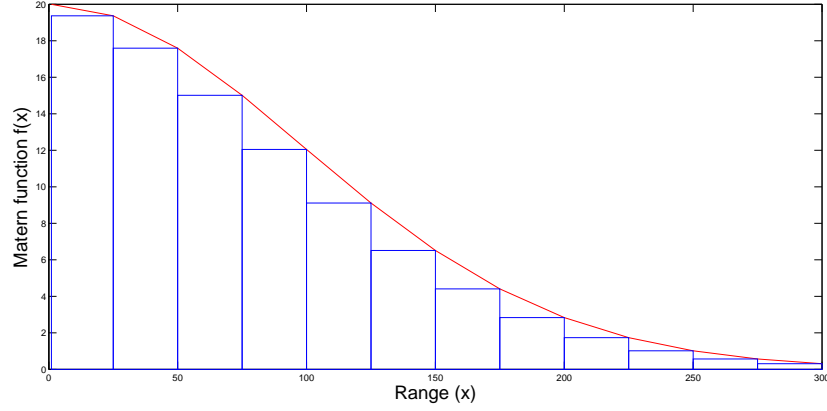


Figure 8.1: The plot gives a visual interpretation of a Riemann sum such that the integral of the function over the range interval $[0, 300]$ is approximated by a sum of rectangles such that the upper right point of the rectangle is the RSS approximated by the Matérn function (in red) at the appropriate range.

With respect to the third bullet (and consistent with the LCM estimation process described in part one), the summation occurs over v bearing windows and this provides a simple way to model the possible anisotropic behavior of the random field in a limited way. The computational goal is now to estimate a series of v random functions at u distinct points. Another way to present this is as a set of vectors.

The vectors are defined with the parameter $\psi_i, i = \{1, \dots, v\}$ that defines an area of coverage (e.g. $\psi_1 = 0$ to 45 degrees and $\cup_{i=1}^v \psi_i = 2\pi$) and a fixed number of increasing ranges $\{r_0, \dots, r_u\}$ (i.e. $r = \{0, 25, 50, \dots, 300\}$) and will be called bearing vectors. They have the following form:

$$\Omega_{nm}^{\psi_i}(x_c) = \{Z_{nm}(\psi_i, r_0), Z_{nm}(\psi_i, r_1), \dots, Z_{nm}(\psi_i, r_u)\} \quad (8.3)$$

A summary of the random field Ω_{nm} can be described by the matrix τ with dimension $v \times u$

$$\tau = \begin{bmatrix} Z_{nm}(\psi_1, r_0) & \dots & Z_{nm}(\psi_1, r_u) \\ \vdots & & \vdots \\ Z_{nm}(\psi_v, r_0) & \dots & Z_{nm}(\psi_v, r_u) \end{bmatrix} \quad (8.4)$$

There are a discrete number of radial estimates around the center point of the LCM $\Omega_{nm}(x_c)$. The number of columns determines the number of increments in radial distance and the number of rows determine the number of increments of the circle (or sphere in \mathbb{R}^3) such that each bearing vector represents the mid-point of each partition.

For example, if there were eight rows in the τ matrix, this corresponds to eight bearing windows such that each window encompassed a 45-degree “wedge”. Suppose a measurement is obtained such that the destination was closest to the random field $\Omega_{nm}^{\psi_i}(x_c)$ centered at position x_c . The bearing of the measurement from source to destination is 30 degrees such that it falls within the bearing window from 0 to 45 degrees ($i = 1$). As will be shown, this measurement is used to improve the estimate of the random function $\mathbf{Z}_{nm}^{\psi_i}$. The range vector is the input into the stochastic random function and the output is the mean (μ) and variance ($var = \sigma^2$) such that $\{\mu_{nm}^{ij}, var_{nm}^{ij}\} = \mathbf{Z}_{nm}^{\psi_i}(\psi_i, r_j)$.

To build the initial GCM during the pre-survey, the following sequence of events is repeated with each message that is promulgated between at least two vehicles in the collaborative system of vehicles :

1. One source vehicle sends a wireless message to another destination vehicle.
2. The destination vehicle records the communication statistics as well as the position of each vehicle (the position is embedded into the wireless message).
3. The measurement is used to improve the estimate of one of the LCMs within the GCM.
4. Selection of which measurement belongs to a particular LCM is based on a minimum distance calculation. In other words, a calculation to determine the centroid closest to the destination vehicle.

It is worth noting the impact of the boundary region in this calculation. There are two types of boundaries an impermeable and permeable boundary. These are represented respectively by the red and yellow boundaries in Figure 8.2. The impermeable boundary refers to the situation when it is not possible to conduct communications beyond this boundary. An of an impermeable boundary for undersea acoustic communications would be a shoreline. In this case it is not possible to communicate beyond this border and the mean estimate to communicate is zero.

A permeable boundary is where it is possible to communicate beyond the boundary but the border exists to define the survey region Ω . An example in Figure 8.2 is the region around the entrance to the Monterey Harbor. Outside this boundary, the mean and variance estimate can still be greater than zero. Impermeable boundaries impact the omnidirectional estimate of communication. Those regions that are near impermeable boundaries will have a poorer assessed ability to communicate prior to any measurements taken when compared to a random field located further from an impermeable boundary.

Any measurements collected within the bearing window are aggregated such that they are considered as part of a realization of a single random function. It is recognized that this is a limiting assumption. It is unrealistic to assume that the anisotropic behavior of a communication connectivity map will adhere to fixed bearing windows. Future research may address the ability to more accurately estimate a random field from a limited number of measurements. One particular technique that may be of utility is known as compressive sensing [46].

The random field can be considered a \mathbb{R}^d dimensional signal. Given an appropriate basis function (such as Fourier or Wavelet basis function) it is possible to use compressive sensing to recreate the signal with minimal errors. Initial research emphasized the requirement for taking random observations. More recently research into adaptive compressive sensing have developed non-random surveying techniques [49], [50], [51].

8.4 Summary

This chapter introduced the Global Connectivity Map (GCM). It extends the Local Connectivity Map (LCM) so that it is possible to estimate the signal strength between any two points within the boundary defined by the map. The GCM is based on the simultaneous estimation of multiple random fields. It does so by partitioning a region into grid cells and



Figure 8.2: The above graphic displays an overlay that represents a visual depiction of a Global Connectivity Map. The red and yellow borders form the permeable and impermeable boundaries for an undersea GCM. The impermeable boundaries are in red and are examples of where no communication is possible outside the boundary. The permeable boundaries in yellow are examples where communication is possible outside the boundary. The gray grid represents the resolution of the map such that the goal of the estimation process to produce an LCM that is centered on the grid cell (in light blue).

each has a signal strength model relative to the center of each cell. Each grid cell produces an overall estimate of the ability to communicate as well as a measure of uncertainty. It also includes a signal strength estimate as a function of distance with regards to several bearing windows emanating outward from the center of the cell.

The next chapter presents a recursive filtering solution for estimating the random functions for each grid cell. This combines together the Kriging estimation techniques in part one with traditional Kalman filtering to produce a Spatial Communications Kalman Filter (SCKF).

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 9:

Building the Global Connectivity Map

With the GCM defined, the next step is to build it. For each grid cell, this requires a procedure for using observations (measurements) to estimate signal strength as a function of range. The procedure should be able to build the map in near real-time as measurements become available and be able to do so in a distributed fashion such that it does not require wireless connections with unlimited bandwidth.

The presented solution combines together Gaussian Process (GP) regression within a recursive Bayes filter. The methodology is similar to GP-Bayes Filters by Ko and Fox [144], [145] but addresses the spatial aspect instead of the temporal domain associated with sequentially estimation of robotic motion through time.

The chapter starts with a discussion of Gaussian Process regression. Next, Gaussian Process Models (GPMs) are introduced. GPMs is a general term that encompasses both regression and classification. It is followed by a discussion of covariance function selection and specifically the Matérn family of functions and the benefits of using a prior for estimating a random function. After that is a review of the discrete Kalman Filter. Then the Spatial Communications Kalman Filter is presented. The chapter concludes with results and comparative analysis.

9.1 Gaussian Process Regression

The use of GP regression permits the estimation of a random function where signal measurements at different ranges and bearings maybe correlated. It is typically a batch procedure—all data is first collected, the procedure is run, and a solution is produced. GP regression is very similar to Kriging. In fact, the equations for the mean and variance are the same. The differences between the techniques are principally associated with the communities that use them. Kriging is a staple of the geostatistics and environmental sciences communities. These communities have placed an emphasis on the semi-variogram. As has been seen, the accuracy of the stochastic model is dependent on the accuracy of the semi-variogram. The weighting function used to produce the Kriging estimator is a spatial relationship to the point of interest and the influence of the spatial relationship is dependent

on the correlation of the measurements. Careful consideration is required in producing an accurate estimator.

GP regression is technique for estimation of a Gaussian random function. The function is completely specified by a mean $\mathbf{E}[f(x)] = \mu(x)$ and covariance $K(x, x') = \mathbf{E}[(f(x) - \mu(x))(f(x') - \mu(x'))^T]$ where the \mathbf{E} is the expectation operator. The random functions under consideration are normally \mathbb{R}^2 and \mathbb{R}^3 but can be generalized to higher dimensions. A Gaussian random function is a collection of random variables of which any finite number have a joint Gaussian distribution.

There are several strengths to the approach. First, it is a non-parametric methodology; the mean surface or curve is not limited by a parametric function's elasticity. Second, the technique provides a measure of uncertainty associated with the mean function. This can be visualized in one and two dimensions as a band around the mean estimate. With reasonable estimates for measurement noise, this can frequently mean that locations with a greater number of measurements will have a narrower uncertainty band with regard to the mean estimate. The ability for estimating a variance is in contrast to other techniques such as Maximum Likelihood Estimation (MLE).

Third, the GP processes include techniques for covariance estimation. For example, we are interested in the impact of a measurement to the surrounding vicinity measured as a range or euclidean distance. In other words the covariance function specifies the covariance between pairs of random variables such that two random variables that are close together will tend to have more similar measurements than random variables that are further apart. Given a covariance function it is possible to learn the function's parameters (these are frequently called hyper-parameters).

9.1.1 The Gaussian Process Model

GPMs are principally used by the Artificial Intelligence (AI) community which (generally) includes Machine Learning and Robotics as sub-disciplines. One of the principle tasks for the Machine Learning community is classification. Common tasks include recognizing handwritten text or image classification where a binary decision is made (yes or no) as to whether an observation of interest belongs to a particular class. With classification, it is no longer possible to use a posterior Gaussian likelihood function since a binary decision is required. One non-linear classification technique is known as logistic regression. This will

be discussed in more detail in the next chapter.

GPMs include an optimization routine to determine the parameters of a selected spatial covariance function. The strength of the approach is that it automates the process of determining the spatial dependence of random variables. Weaknesses in the technique is the optimization process requires all the data in advance, may take significant processing time and may not converge to an optimal solution. In previous chapter there has been an emphasis on Kriging. The reason for the switch from Kriging to GPMs is to remain consistent with recent similar robotics research [144], [145] and to ensure that it is not confused with a different approach—called the Kriging Kalman Filter [146]. The building of the LCM could be done with GPMs but it was felt that structural analysis plays a critical role in building the GCM and understanding the stochastic regression model is incomplete without it.

The Gaussian Process Model starts with a data set $D = (\mathbf{X}, \mathbf{z})$ where \mathbf{X} represents known positions ($\mathbf{X} = \{x_1, x_2, \dots, x_n\}$ and x_i is of dimension d) and \mathbf{z} is a set of measurements ($\mathbf{z} = \{z_1, z_2, \dots, z_n\}$) such that z_i is a measurement that corresponds to the position x_i . It is assumed that there is measurement noise associated with the observational process:

$$z_i = f(\mathbf{x}_i) + \varepsilon \quad (9.1)$$

where a measurement is equal to a function plus an error term. The function has as input a vector of known points \mathbf{x}_i and ε is the error term that is assumed to have an independent identically distributed (i.i.d.) Gaussian distribution with zero mean and covariance K (i.e. $\mathbb{N}(0, K)$).

Of interest is the estimation of a signal measurement at a point of interest (POI) (x_0). The GPM defines the predictive mean distribution given the test data (D) as:

$$GP_\mu(x_0, D) = \mathbf{k}^T (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{z} \quad (9.2)$$

where \mathbf{k} is short for $\mathbf{k}(x_0, \mathbf{x})$ which is the covariance between the point of interest (x_0) and the selected data points (\mathbf{x}). The selected data points are a neighborhood of points around the POI that are used to determine the estimate. the matrix \mathbf{K} is short for $\mathbf{K}(\mathbf{x}, \mathbf{x})$ and is the covariance between the test points. Notice that except for the noise term multiplied by the identity matrix (σ^2), the formula for the predictive mean estimate is the same as the mean estimate for ordinary Kriging.

The Gaussian Process variance calculation is given by

$$GP_{\Sigma}(\mathbf{x}_0, D) = k(x_0, x_0) - \mathbf{k}^T (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{k} \quad (9.3)$$

The term $k(x_0, x_0)$, represents the self-covariance of the POI. In the Kriging terminology this was defined as the nugget. Derivation of these equations are available in several resources including [62], [68]. Again the GPM predictive variance is almost the same as the calculation for the ordinary Kriging variance with the exception of the $\sigma^2 \mathbf{I}$ term which is an additional noise or error term.

Generally considered part of GPM is the learning of hyper-parameters associated with the parametric covariance function. By looking at a batch of observation measurements it is possible to maximize the log marginal likelihood of the training data. Mathematically,

$$\theta_{max} = \operatorname{argmax}_{\theta} \{\log(p(\mathbf{z}|X, \theta))\} \quad (9.4)$$

where the log term is given by

$$\begin{aligned} \log(p(\mathbf{z}|X, \theta)) = & -\frac{1}{2} \mathbf{z}^T (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} \mathbf{z} \\ & -\frac{1}{2} \log |\mathbf{K} + \sigma^2 \mathbf{I}| - \frac{n}{2} \log 2\pi \end{aligned} \quad (9.5)$$

and n is the number of observations. It is common to use a conjugate gradient optimization routine to determine appropriate values for the hyper-parameters [68]. To do so requires the partial derivatives with respect to the hyper-parameters. Note, that there is no guarantee that the process will produce a single optimal result since the objective function may have several optimal solutions. In general, the procedure is known to work well. This means that given a candidate function with only a few parameters for characterizing the function, the procedure converges to a reasonable value within a reasonable time frame. As described above, of critical importance to the accuracy of the model is the selection of an appropriate covariance function.

9.1.2 The Matérn Covariance Function

The Matérn family of covariance (or correlation) functions is considered an appropriate function for spatial processes [15]. In part one, the function was introduced as a candidate function for fitting the empirical semi-variogram data. In this case, it will be used as a covariance function to determine the influence of new measurements on the random function in terms of the points located in the neighborhood of the measurement. The equation is given by:

$$\rho(s; \alpha, \beta, \nu) = \alpha \left[\frac{1}{2^{\nu-1} \Gamma(\nu)} \left(\frac{s}{\beta} \right) H_{\nu} \left(\frac{s}{\beta} \right) \right] \quad (9.6)$$

where s is the distance metric between observations, H_{ν} is a modified Bessel function of the third kind of order ν and Γ is the Gamma function. There are three function parameters: α scales the function at the y-axis intercept as the distance (s) approaches zero, β determines how quickly the function degrades over increasing distance and ν is a smoothness parameter ($\nu > 0$) [147]. The utility of the function comes from the fact that selection of the different parameters result in a variety of curve shapes that are continuously differentiable [20] and varying values of ν result in a more gradual approach to zero as distance increases. A selection of curves are displayed with the function parameters in Figure 9.1.

In summary, the Matérn function is a reasonable selection for a covariance function. From part one, it was shown that the covariance at a lag equal to zero is equivalent to the semi-

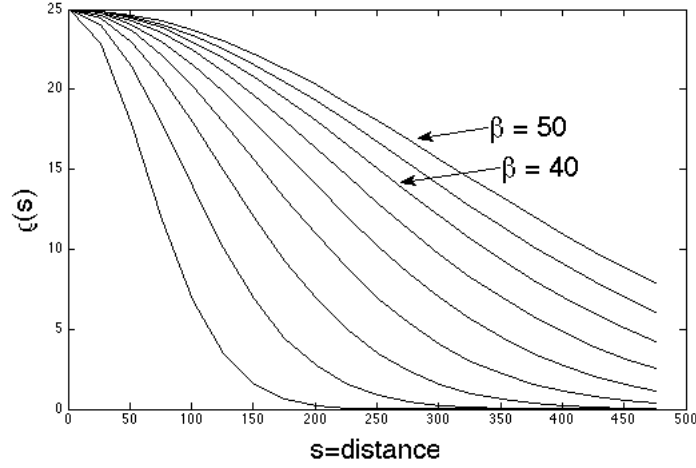


Figure 9.1: A selection of Matérn covariance functions.

variogram sill. One must be careful with the Matérn function with evaluations at zero ($s = 0$) since the function is undefined. Also the semi-variogram range defines the lag distance at which the variables are independent. This is the point where the covariance should be zero.

9.1.3 Applying Gaussian Process Modeling to RSS Data

This section shows results using a GPM approach for estimation of a random function for the estimation of RSS. Figure 9.2 shows a collection of RSS measurements in blue. These measurements were taken relative to a static undersea acoustic communications buoy located at position S_1 . A subset of the measurements were selected for calculating the random function. It included 122 measurements in a bearing window of between 165 and 195 degrees.

Figures 9.3 and 9.4 show the results of applying the GPM to 50 and 122 of the measurements. Notice there were no measurements within approximately 50 meters of the static buoy. With no data to improve the estimate the RSS random function has a mean and variance estimate from 0 to 25 meters which remains unchanged from the *a priori assumption* (the example started with a uniform prior of $\mathcal{N}(10, 9)$). For the random function distances from 50 to 300 meters, the measurements show a relatively small drop off in RSS over increasing distance and the variance has been significantly reduced from the original assumption.



Figure 9.2: The graphic depicts the locations of acoustic signal strength measurements (in blue) from Monterey Harbor and the boundaries of the bearing window (in red) from 165 to 195 degrees where the points within the boundaries are used to estimate the random function.

Given shape of the rest of the random function it is improbable to think that the RSS near the static acoustic modem is significantly less than the RSS at 100 meters from the acoustic modem. Figures 9.5 and 9.6 show the random function as a result of the 50 and 122 measurements, respectively, but these estimates use an Matérn function ($\alpha = 20, \nu = 20, \beta = 20$) as the prior mean. Note that the prior is a Matérn function modeling the RSS as a function of range, this should not be confused with the Matérn covariance function that is used to establish the relationships between measurements based on distance.

9.2 Optimal Estimation

This chapter has discussed the use of GPMs for estimating a random function. Of interest is the ability to estimate several random functions relative to bearing windows emanating out from the center point of a grid cell. The collection of cells make up the GCM. Depending on the desired resolution of the GCM there may be a large number of concurrent estimation problems.

GPMs provide good results but are not the ideal tool for this particular problem for two reasons: First, we are interested in building the GCM in near real-time. GPMs are designed

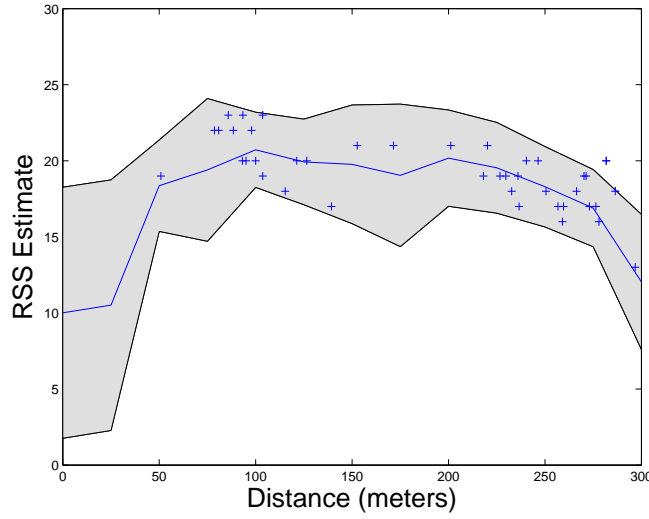


Figure 9.3: The plot displays an estimate of a random function based on 50 S_1 measurements in the bearing window from 165 to 195 degrees.

to use a batch of observations to generate random function estimates. Since measurements will generally arrive to vehicles incrementally it takes time before enough measurements are collected before a batch run is conducted. This latency is undesirable since one would like to build the map as the measurements become available.

Second, it would be necessary to send the raw measurements through the network so that other vehicles can each run batch updates. This is inefficient and could potentially negatively impact bandwidth restricted networks. A preferred solution would be a recursive filter that provided improved estimates as each measurement became available. Additionally, the ability of a filter to improve an estimate through a Markov assumption potentially alleviates overburdening the wireless network.

A potentially better solution is to combine aspects of the GPM with a recursive optimal estimation filter. The GPM provides the ability to use measurements to improve the random field mean and variance around a neighborhood. Where the neighborhood is defined by the spatial covariance between random variables. The optimal estimation filter provides the ability to recursively improve the random field estimate through individual (or batch) measurements. Secondly, the mean and variance for a particular partition can be sent throughout the wireless network such that vehicle receiving the information can im-

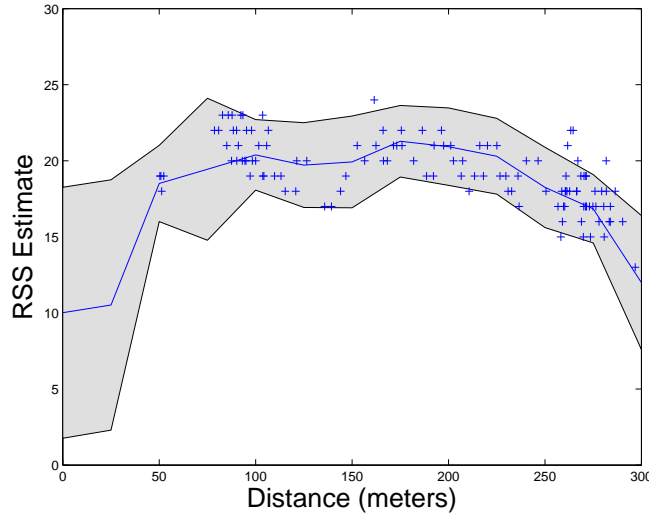


Figure 9.4: The plot displays an estimate of a random function based on 122 S_1 measurements in the bearing window from 165 to 195 degrees.

prove GCM resolution with less bandwidth requirements.

Normally, a recursive optimal estimation filter is associated with a temporal problem like the position of an unmanned system over time. In this application we are interested in the optimal estimation for the spatial domain. Techniques for estimation of a random field frequently assume that the data set is already available. Due to the nature of this collaborative communications mapping problem we are interested in a technique for the recursive on-line solution in the spatial domain.

9.2.1 Kalman Filter

The Kalman filter is a Bayesian recursive filter that is an optimal estimator of a linear system [148]. An optimal estimator “... is a computation algorithm that processes measurements to deduce a minimum error estimate of the state of a system by utilizing knowledge of system and measurement dynamics, assumed statistics of system noises and measurement dynamics, assumed statistics of system noises and measurement errors, and initial conditional information.” [149]

The Kalman filter is a popular choice for temporally-based estimation problems for its relative simplicity and its ability to minimize estimation error in a well-defined statistic

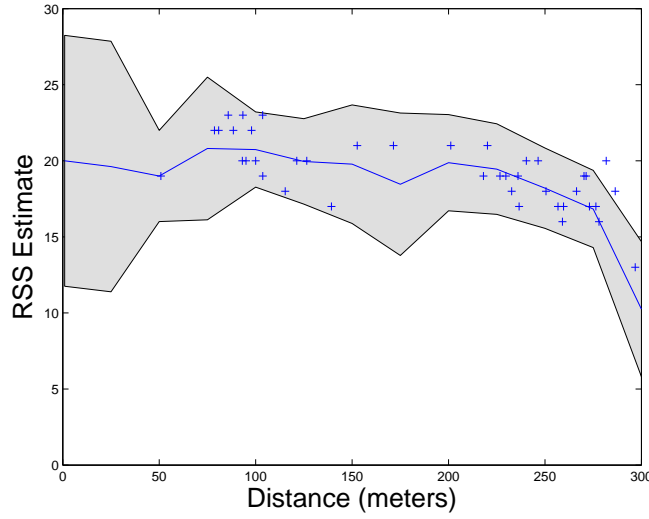


Figure 9.5: The plot displays an estimate of a random function based on 50 S_1 measurements in the bearing window from 165 to 195 degrees using a Matérn prior mean function.

manner. Conversely, the disadvantages include the computational burden associated with matrix inversion and sensitivity to a-priori models.

The discrete, standard linear Kalman filter is a recursive algorithm that has inputs and outputs. The inputs are the state vector $\mathbf{x}_{k|k}$, covariance matrix $\mathbf{P}_{k|k}$ and measurements \mathbf{z}_k ¹⁰. The output of the filter is a new predictive estimate for the mean $\mathbf{x}_{k+1|k+1}$ and covariance $\mathbf{P}_{k+1|k+1}$. This output is then used together with new measurements for the next iteration of the filter.

The filter is defined by a system model $\mathbf{x}_{k+1|k} = \Phi \mathbf{x}_{k|k} + \mathbf{w}_k$ where Φ is a transition matrix and \mathbf{w}_k is a system error term that is assumed to have a Gaussian distribution with mean zero and an error covariance matrix \mathbf{Q} (otherwise written $\mathbf{w}_k \sim \mathcal{N}(0, \mathbf{Q})$).

The measurement model $\mathbf{z}_k = \mathbf{H} \mathbf{x}_{k|k} + \mathbf{v}_k$, where \mathbf{z}_k are the measurements, \mathbf{H} is a transition matrix relating the state vector to the measurements and \mathbf{v} is the measurement error with mean zero and an error covariance matrix \mathbf{R} ($\mathbf{v}_k \sim \mathcal{N}(0, \mathbf{R})$). The steps of the recursive

¹⁰The notation convention used for the Kalman reflects the Bayesian estimate of the mean and covariance at different stages of the algorithm. $\mathbf{x}_{k|k}$ reflects the current state input mean estimate, $\mathbf{x}_{k+1|k}$ is the update of the state and $\mathbf{x}_{k+1|k+1}$ is the prediction of the mean state vector. The same subscript notation holds for the covariance matrix \mathbf{P} .

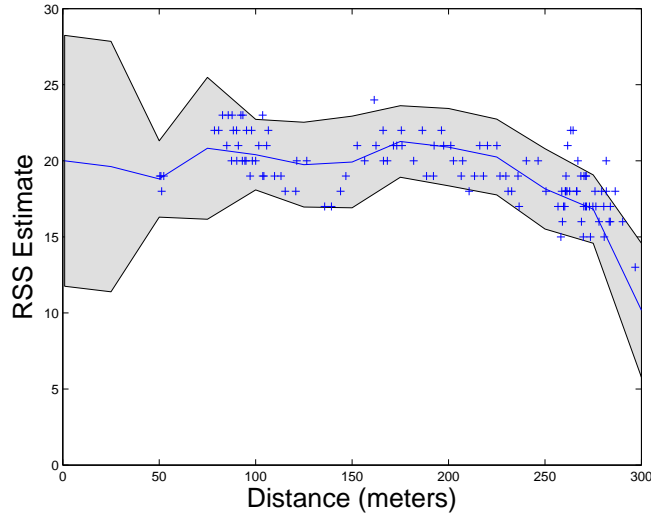


Figure 9.6: The plot shows the estimate of a random function based on 122 S_1 measurements in the bearing window from 165 to 195 degrees using a Matérn prior mean function.

algorithm consists of the following:

1. Algorithm: Linear Kalman Filter (LKF)
2. Input: LKF($\mathbf{x}_{k|k}$, $\mathbf{P}_{k|k}$, \mathbf{z}_k)
3. $\mathbf{x}_{k+1|k} = \Phi_k \mathbf{x}_{k|k}$
4. $\mathbf{P}_{k+1|k} = \Phi_k^T \mathbf{P}_{k|k} \Phi_k + \mathbf{Q}_k$
5. $\mathbf{G}_k = \mathbf{P}_{k+1|k} \mathbf{H}_k^T (\mathbf{H}_k \mathbf{P}_{k+1|k} \mathbf{H}_k^T + \mathbf{R}_k)^{-1}$
6. $\mathbf{x}_{k+1|k+1} = \mathbf{x}_{k+1|k} + \mathbf{G}_k (\mathbf{z}_k - \mathbf{H}_k \mathbf{x}_{k+1|k})$
7. $\mathbf{P}_{k+1|k+1} = (\mathbf{I} - \mathbf{G}_k \mathbf{H}_k) \mathbf{P}_{k+1|k}$
8. Output: $\mathbf{x}_{k+1|k+1}$, $\mathbf{P}_{k+1|k+1}$

Step five of the algorithm is the calculation of the Kalman gains matrix \mathbf{G} . This matrix is then used in step six and seven to calculate the prediction of the mean and system covariance.

9.2.2 Combining the KF with GPM

Of interest is the RSS estimation at several discrete range values. For a single bearing window, the state vector for the recursive filter are RSS estimates at fixed range positions relative to the central cell node. There is a filter running for each bearing window such that

the collection of bearing windows provide an anisotropic estimate of the random field.

One is not guaranteed that measurements will be observed at exactly these ranges. For this reason, it is necessary to approximate mean and variances based on a covariance relation between random variables. This will permit a measurement to improve the state vector elements where the covariance between the location of the measurement (\mathbf{x}_*) and the state vector (\mathbf{x}) is non-zero. The filter is called the Spatial Communications Kalman Filter (SCKF). The name reflects the spatial and recursive components of the filter. It consists of the following steps:

1. Algorithm: SCKF
2. Input: SCKF($\mathbf{x}_{k|k}, \mathbf{P}_{k|k}, (\mathbf{x}_*, \mathbf{z}_k)$)
3. $\mathbf{x}_{k+1|k} = \Phi_k \mathbf{x}_{k|k}$
4. $\mathbf{P}_{k+1|k} = \Phi_k^T \mathbf{P}_{k|k} \Phi_k + \mathbf{Q}_k$
5. $\hat{\mathbf{z}}_k = f(\mathbf{x}_{k|k}, r)$
6. $\text{diag}(h) = \mathbf{z}_k / \mathbf{x}_{k+1|k}$
7. $\mathbf{H} = sf(\text{diag}(h), \mathbf{K})$
8. $\mathbf{x}_{k+1|k+1} = \mathbf{x}_{k+1|k} + GP(\hat{\mathbf{z}}_k, \mathbf{z}_k)$
9. $\mathbf{G}_k = \mathbf{P}_{k+1|k} \mathbf{H}_k^T (\mathbf{H}_k \mathbf{P}_{k+1|k} \mathbf{H}_k^T + \mathbf{R}_k)^{-1}$
10. $\mathbf{P}_{k+1|k+1} = (\mathbf{I} - \mathbf{G}_k \mathbf{H}_k) \mathbf{P}_{k+1|k}$
11. Output: $\mathbf{x}_{k+1|k+1}, \mathbf{P}_{k+1|k+1}$

The difference between the standard discrete Kalman Filter and the SCKF begins with step five. A measurement estimate ($\hat{\mathbf{z}}_k$) is obtained using a function ($f(\cdot)$) which has, as an input, the current state vector mean estimate ($\mathbf{x}_{k|k}$) and the range (r) of the current measurement from the signal source (equivalently the grid center of a LCM).

The \mathbf{H} matrix is calculated with steps 6 and 7. Step six is an element-wise division of the measurement estimate ($\hat{\mathbf{z}}_k$) by the state vector update $\mathbf{x}_{k+1|k}$. Next in step seven, a modified step function (sf) is employed such that for all diagonal elements ($i = 1 \dots n$).

$$H_{ii} = \begin{cases} 0 & \text{if } K_{ii} = 0 \\ h_{ii} & \text{otherwise} \end{cases}$$

This ensures the \mathbf{H} matrix properly scales the relation between the state vector and mea-

surement by considering the spatial influence of the covariance function.

Finally the prediction of the state vector $\mathbf{x}_{k+1|k+1}$ is modified from a standard Kalman filter. It consists of the state vector update $\mathbf{x}_{k+1|k}$ plus an innovation that is a result of the application of the Gaussian Process Model. Normally the predictive mean equation is given by equation (6), Now we are interested in the mean of the innovation and it is given by:

$$GP(\hat{\mathbf{z}}_k, \mathbf{z}_k) = \mathbf{k}^T (\mathbf{K} + \sigma^2 \mathbf{I})^{-1} (\hat{\mathbf{z}}_k - \mathbf{z}_k) \quad (9.7)$$

The result of the GP function is a column vector with the same dimension of the state vector. In other words, in many cases, the measurement will be a single observation and using the assumed covariance relationship the GP function determines an innovative estimate where $cov(\mathbf{x}_*, \mathbf{x}_i) > 0$, $\mathbf{x}_i \in \mathbf{x}$.

An important consideration is the observability of the system. It defines one's ability to determine the state vector from the measurements. It is calculated using the following matrix Ξ [149]:

$$\Xi = \begin{bmatrix} \mathbf{H}^T | \Phi \mathbf{H}^T | \dots | (\Phi^T)^{n-1} \mathbf{H}^T \end{bmatrix} \quad (9.8)$$

where n is the length of the measurement vector. As long as the matrix (Ξ) has a rank greater than n the system is observable. Given that the transition matrix Φ is an identity matrix, Ξ is observable as long as the multiplicative combination of \mathbf{z} and \mathbf{x}_k^{-1} has a rank of n . This implies that neither vector elements of \mathbf{z} or \mathbf{x}_k^{-1} can be zero, there is a positive definite covariance function \mathbf{K} and measurements are collected along all the ranges of the state vector.

9.2.3 An SCKF Example

In an acoustic modem dataset, Figure 9.7 shows raw RSS measurements taken relative to a static buoy located at local image coordinates (382,312) and global coordinates (Lat/Lon

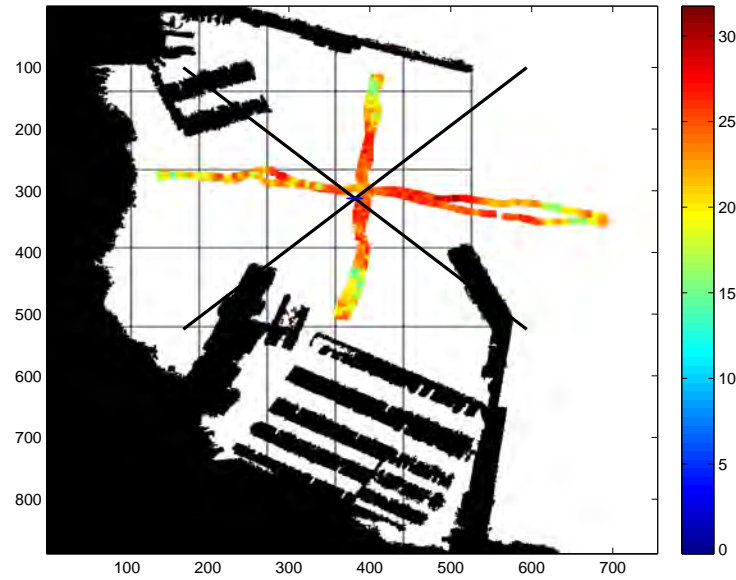


Figure 9.7: The figure show a dataset of RSS measurements taken in March 2013 in Monterey Harbor, CA. It is partitioned into quadrants north, east, south and west

36.607352 , -121.891018). The data is partitioned into four quadrants—north, south, east and west as is shown in Figure 9.7. With each quadrant, the goal is to estimate the random function using the SCKF filter. Figure 9.8 shows a plot of the four RSS measurements as a function of range.

Figure 9.9, shows the result of a sequential running of the SCKF such that one RSS measurement was used for estimating the system state vector at each time step. It includes the mean of the state vector, the system variance and the raw RSS measurements. The mean of the random function is drawn as a continuous line by connecting the mean estimates and this is done with the system variance as well. Figure 9.10 shows the convergence of the system by plotting the trace of the covariance matrix. All elements of the covariance matrix P were initialized to 10. The convergence of the covariance is rapid since it is in effect using an a priori assumption about the measurements in step five of the SCKF algorithm.

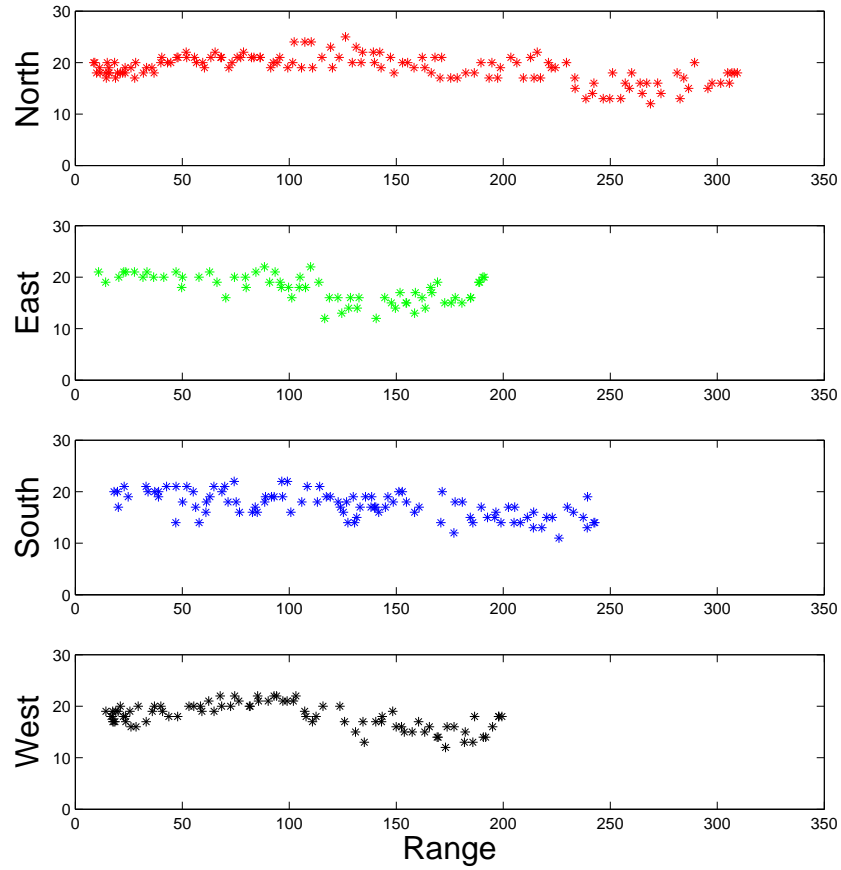


Figure 9.8: The figure shows the north, south, east and west plots of RSS values over increasing range. This is the same dataset the previous figure.

9.2.4 Cubic Spline Interpolation

An important consideration in this methodology is the selection of a function for representing the RSS state vector. Key to the implementation is a function which has the elasticity for taking into consideration a measurement which is significantly different from previous estimates. This is best illustrated through an example.

One can start with a reasonable assumption with regards to the communications channel. This is reflected in the filtering process in two ways:

1. The *a-priori* initialization of the state vector (\mathbf{x}_0) and covariance matrix (\mathbf{P}_0).

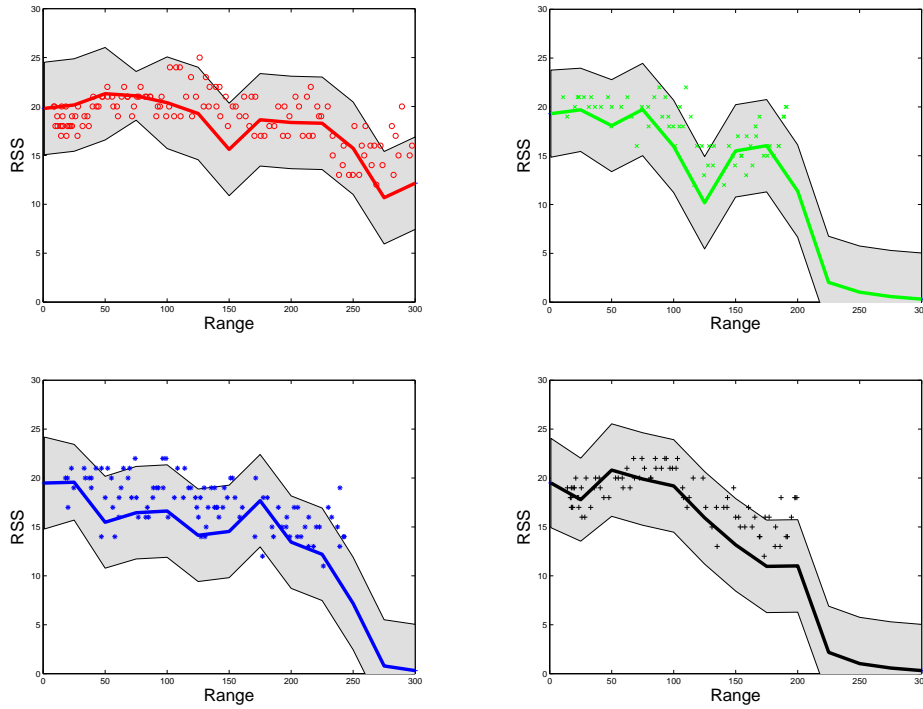


Figure 9.9: *SCKF Results on the North, East, South and West data sets.*

2. The function used to approximate the random function in step five of the filter. ($\hat{\mathbf{z}}_k = f(x_{k|k}, \mathbf{z}_k, r)$)

Unless a function has the flexibility in its curvature to reflect innovation in the filtering process, the resulting mean and covariance estimate suffer as a result. The blue line in Figure 9.11a shows an initialization of the RSS mean function with the Matérn function with parameters (20,16,20). The red star is a new measurement and the red line shows the change in the function as a result of the new measurement. In this example, a single measurement impacts the mean estimate for almost all elements of the mean state vector \mathbf{x}_i .

Figure 9.11b shows the same sequence of events but instead of a Matérn function, a cubic spline function is used. Cubic spline interpolation is a common technique for creating a piecewise continuous curve [150]. It is characterized by a cubic polynomial for each defined interval such that each interval has its own set of coefficients. The curve is made smooth by ensuring the the first and second derivatives at the connecting intervals are con-

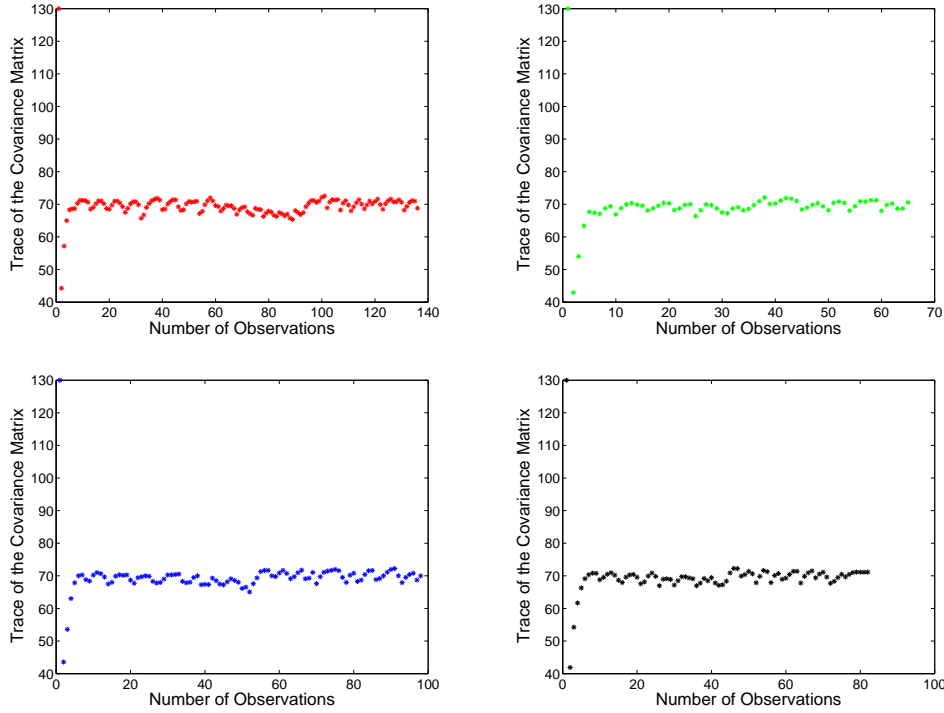


Figure 9.10: *The convergence of the system covariance. The four plots show the trace of the SCKF covariance matrix for the North (Red), East (Green), South (Blue) and West (Black) data sets.*

tinuous.

Figure 9.11b shows the result of the change to the curve using the cubic spline interpolation when a new RSS measurement is used to update the mean estimate. The changes are “local”, meaning they are fairly consistent with the covariance influence of the measurement while the changes to the Matérn function were not.

9.2.5 SCKF and GPM comparison

It is possible to now compare the SCKF with the GPM. Figure 9.12, shows the results of using the Gaussian Process Model (or Kriging) on the same data set. When comparing Figures 9.9 (SCKF results) and 9.12 (GPM results), one can notice several trends. First GPM estimates in the East, South and West datasets do not show the drop off in the 200–300 meter range that the SCKF results show. The GPM use an optimization routine for estimating the Matérn mean function hyper-parameters and the combination of a parametric mean function together with no data available within this area means that the data in the 0–200

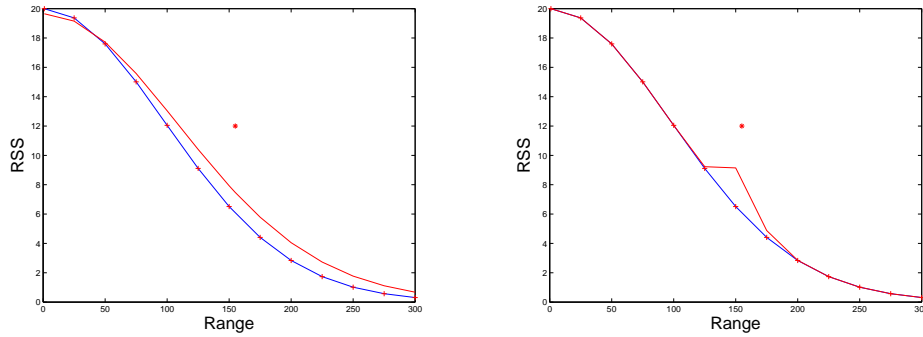


Figure 9.11: On the left is a Matérn function and on the right is a cubic spline. The Matérn function impacts all estimates along the range while the cubic spline has minimal impact to the surround points.

meter range has the impact of increasing the mean estimate at the tail end. Conversely, with the SCKF, the lack of data in the 200–300 meter range means that the *a-priori* assumption still holds.

Second, the GPM use the entire dataset to produce the mean and variance estimate. This results in the mean and variance to nicely represent the data in that it is centered on the data. Conversely, the SCKF mean estimate (especially in the East, South and West) tends to slightly underestimate the true mean and variance of the entire data set. Since this is a recursive algorithm that starts with a conservative *a-priori* estimate, this is to be expected.

9.3 Summary

This chapter has focused on the building of the Global Connectivity Map. The GCM is defined, in part, by a discrete number of random functions anchored by a common central point. Given the proper measurements, the random function estimation for each GCM tile (or cell) permits one to estimate the potential anisotropic behavior of the field. A Spatial Communication Kalman Filter (SCKF) was developed that combines together a GPM and Linear Kalman Filter for the near real-time estimating of a random function. The key augmentation to the Kalman Filter is the use of a function for calculating a measurement estimate. A cubic spline was selected as an appropriate function since it minimally impacts the estimate in the neighborhood surrounding the new measurement. The next chapter considers the compendium of LCMs as a network and looks to simplify the GCM by looking for evidence of similarities between LCMs.

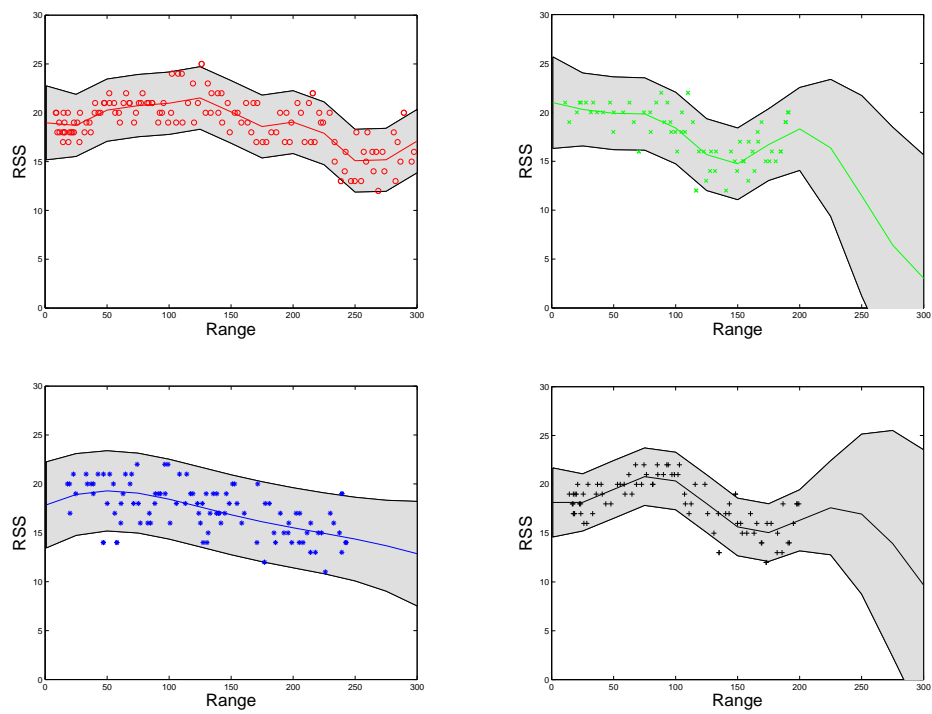


Figure 9.12: *Gaussian Process Model results on the North (Red), East (Green), South (Blue) and West (Black) datasets.*

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 10:

Refining the Map

In the previous chapter, a Spatial Communication Kalman Filter (SCKF) was introduced for the building of the Global Connectivity Map (GCM). A natural goal in building the GCM is to minimize the number of measurements necessary to create the GCM. In the last chapter, strategies were described to reduce the necessary measurements. This included starting with *a priori* assumptions of how the RSS signal degraded over increasing range.

For this chapter the principal topic is understanding the relation between grid cells in the GCM. The purpose is two-fold. First, if two areas have similar signal characteristics it may be possible to reduce the number of measurements by combining the grid cells together such that a single model can be used to represent both regions. Another way to think about this is that an arbitrary grid was defined as an overlay on the survey region. Instead, it would be useful to define the grid cell boundaries by the uniformity of the signal characteristics. In this respect, the goal is a transformation from an initial grid representation to one which considers the environmental impact of communication.

A more lofty and long range goal is the ability to reveal aspects of the environment that impose constraints on communication. An example in the underwater domain would be the identification of a muddy ocean floor which absorbs acoustic energy and dramatically decreases network connectivity. This might be considered a second order sensing system where the modem is used primarily for wireless communications but a secondary byproduct of the analysis of signal statistics is that they can help identify a phenomena (the muddy ocean bottom) that impacts communication. In other words, through model development and signal measurements one might be able to uncover a hidden model variable which has explanatory power.

The chapter begins with the problem statement. It follows with a discussion of recent, relevant research. Next, a brief introduction of the mathematical and networking framework that is used to address the problem. That is followed by a simulated example and the chapter concludes with a discussion of the convergence properties of the technique.

10.1 Problem Statement

Of interest is the ability to understand the structure of a GCM. Specifically of interest, is the ability to determine the dependency between neighboring Local Connectivity Maps that make up the Global Connectivity Map. The problem can be posed as a network structure learning problem. Represent the GCM as a graph $(G = (V, E))$ where each vertex (V) is a one-dimensional model associated with a particular bearing window of a LCM and each edge (E) represents a correlation or dependency relation between the models.

The presence of an edge represents the case where a measurement received in one grid cell validates in the model of the connected vertex. Validation is the process where a received measurement “fits” within a specified tolerance of the neighboring model. The lack of an edge represents the opposite—a lack of dependency between the nodes such that measurements from one node did not validate in the model of the connected node. The general problem statement is how to estimate a time-varying network using a graph structure from n independently distributed measurements.

10.2 Recent Research

This work builds on the emerging theory on temporal networks [151]. Over a great variety of systems in society, engineering and science a common theme is the attempt to gain greater understanding of networks. The goal is, in part, driven by the influence of the internet over recent generations. The temporal aspect of networks is to determine the nature of how network connections change over time.

For this chapter, the goal is the development of a methodology for learning how regions of a wireless communication networks are similar. It leverages recent work in time-varying network analysis. It is worth noting that the emphasis for this work is not on the analysis of how the network changes over time (as it can be with time-varying networks) but on the identification of conditional dependencies between neighboring graph nodes. That said, the ability to analyze the environmental impact on communications through temporal analysis could be useful. This is left for future work.

In this case, the network structural estimation process is made more difficult for two reasons. First, there is little initial information available about the environment and second, the process of network connectivity discovery must be on a near real-time basis. This chapter

seeks to extend existing techniques to address both of these issues.

To address the first issue, the previous chapter developed a novel filtering solution for building the LCM model—the Spatial Communications Kalman Filter. As measurements become available they are used for refining estimates of the communications channel. This filter will now be used for the model validation process of the learning network structure portion.

The research area of central importance to the chapter is learning network structure. Of particular interest is a technique using pairwise Markov Random Fields (MRF) with ℓ_1 -regularized logistic regression. There are a series of several papers that are influential. Wainwright et al. presented an ℓ_1 -regularized logistic regression model where the key to learning network structure was to learning the neighborhood structure around each vertex. Another paper by Ravikumar, Wainwright and Lafferty delve into the technique further and provide bounds and computational complexity for the number of samples as a function of maximum neighborhood size and the number of nodes [152].

More recently Kolar, Xing and Le Song have extended the technique for estimating time-varying networks [153], [154], [155] and [156]. Here the difference is how the network changes over time. Examples include voting records of the US Senate and the interactions of genes of an insect over different stages in the lifecycle. The emphasis is on a dataset that changes smoothly over relatively small time changes. The data is weighted with a temporal non-negative kernel so that more recent network interaction is weighted more heavily than older information.

10.3 Overview

Similar to the previous chapter, this work leverages research in a branch of Artificial Intelligence known as Machine Learning. In the last chapter, the focus was on linear regression using Gaussian Process Models. In this chapter the focus is on classification. Informally this is the process by which a computer is tasked with adapting to new circumstances, making decisions or classifications from inputted data.

A common example of a classification problem is face recognition. The goal is to correctly classify a name with a person's photograph. The input might be a picture of a person that is loaded into the computer and the output is a binary decision to correctly classify the

picture as being a particular person (or not). A classic approach is to train a classifier—given examples of the input and the output—a function is developed that can be then used with new data for classification.

This is an example of supervised learning—learning a function from inputs and outputs. There are several other types of machine learning—unsupervised, reinforcement and semi-supervised. Unsupervised learning involves learning the function when no correct output is provided and reinforcement learning is the most general case; it involves inferring behavior through learning how the environment works.

Semi-supervised learning falls between supervised and unsupervised learning where there is both labeled and unlabeled data. Techniques in semi-supervised learning seek to use this combination of data types such that the resulting performance is better than what would be possible with either supervised or unsupervised learning.

Of interest for this chapter is a variant of semi-supervised learning. In this case, there is a temporal sequence of measurements and validation model is used for estimating correlation between vertices. The Probabilistic Graphical Model's (PGM) performance is coupled with the validation model. The validation starts with little information but improves as new measurements are provided. The methodology of learning network structure is dependent on the validation model, so it too will converge as the validation model converges.

The goal is to simultaneously learning the random functions of all LCMs while learning the network structure between neighboring GCM lattice cells. To accomplish this, the results of the last chapter for near real-time signal characterization are combined with learning network structure in a PGM. Normally, learning network structure is accomplished through a classification procedure—there is an existing database of previously collected information (known as the design matrix in the linear regression community) and a model is constructed which produces the most likely candidate network that supports the data. A common technique is the use of optimization techniques such as Maximum Likelihood Estimator over a particular network configuration.

The structure of this problem is a part of a larger set of tasks where one is interested in understanding the dependency or correlation between entities but there is no direct measurement of the underlying relationship. Instead there is a time series of measurements

available at each entity. Other examples of related tasks including ascertaining the network structure of legislators based on voting rights [157], looking at the time varying network of gene regulation in biologic systems [154].

This chapter draws from recent research into learning structure from graphical models. The core idea is to start with a fully-connected graph where each vertex represents a lattice cell. Each of the cells start with an initial LCM model that has a large degree of uncertainty. As measurements become available the LCM estimate for a lattice cell starts to improve. All measurements obtained at a particular lattice cell are validated each neighbor's model. If it does validate—if the measurement is within the bounds of the current neighbor's model—the link (or factor) between the nodes is strengthened, if not, it is weakened. Periodically, the network is re-evaluated to assess the changes in the structure. This is accomplished using an optimization technique known as ℓ_1 regularized logistic regression. The result provides an on going assessment of the GCM network structure as well as the latest estimate of network structure.

The chapter begins with review of recent research. Next, PGM's are introduced with an emphasis on Pairwise Markov Random Networks. It is followed by a discussion of logistic regression and ℓ_1 regularization. It finishes with simulation results. What differentiates this approach from previous research is the use of the improving estimate of the random functions to strength or weaken estimates of the network structure. The chapter finishes with a simple example of the technique based on simulated acoustic modem signal statistics data.

10.3.1 Probabilistic Graphical Models

PGMs "... use a graph-based representation as the basis for compactly encoding a complex distribution over a high-dimensional space" [71]. The primary benefits of PGMs are they provide a visual description of a probabilistic model. Frequently through visual inspection of the PGM, one is able to determine and understand the conditional dependence relationships in the model. Second, the computation required to perform inference and learning are expressed through techniques in algebraic graph theory which nicely represent underlying mathematical expressions.

A PGM is a graph (G) consists of nodes or vertices (V) and edges (E) such that $G = (V, E)$. The nodes are random variables and the edges are probabilistic relationships between the

variables. There are generally two types of graphical representations, Bayesian and Markov networks. They differ in the types of edges contained in the graph. A Bayesian network uses a directed graph where the directional edges go from a source to a target. A Markov network (equivalently known as a Markov Random Field (MRF)) features undirected edges between nodes such that it is possible to traverse in either direction. It is also noteworthy, that PGMs can be represented by Factor graphs where in addition to nodes and edges there are factors which represent relationships between nodes.

10.3.2 Markov Network Example

We are principally interested in Markov networks since correlations between random fields are bi-directional. A simple example helps to illustrate how a Markov network is constructed and how inference is accomplished. Figure 10.1, shows a simple connected, undirected graph with three nodes and three edges. The edges are defined by factors (ϕ_1, ϕ_2, ϕ_3) such that the factors represent relationships between the nodes (e.g., $\phi_1(A, C)$ represents the relationship between node A and C). In the example there are two possibilities between the nodes—the possibility the nodes are similar or dissimilar, this is represented by $\phi_1(A, C) = \{a^1 c^1, a^0 c^0\}$ respectively. These relationships are assigned numerical values that don't necessarily need to sum to one.

The table of the right side of Figure (10.1), lists the discrete joint probability space which is a combination of all the possible assignments to the variables. The table gives assignments to each of the variables such that it is now possible to determine the normalized probability for each of the discrete hypotheses. This is calculated by:

$$P(a, b, c) = \frac{1}{Z} \phi_1(a, b) \cdot \phi_2(b, c) \cdot \phi_3(a, c)$$

$$Z = \sum_{a, b, c} \phi_1(a, b) \cdot \phi_2(b, c) \cdot \phi_3(a, c) \quad (10.1)$$

where Z is the normalizing constant known as the partition function. It takes the factors, which summarize the relationships between variables, and produces probability assignments that sum to one.

Markov Random Field 3 node example

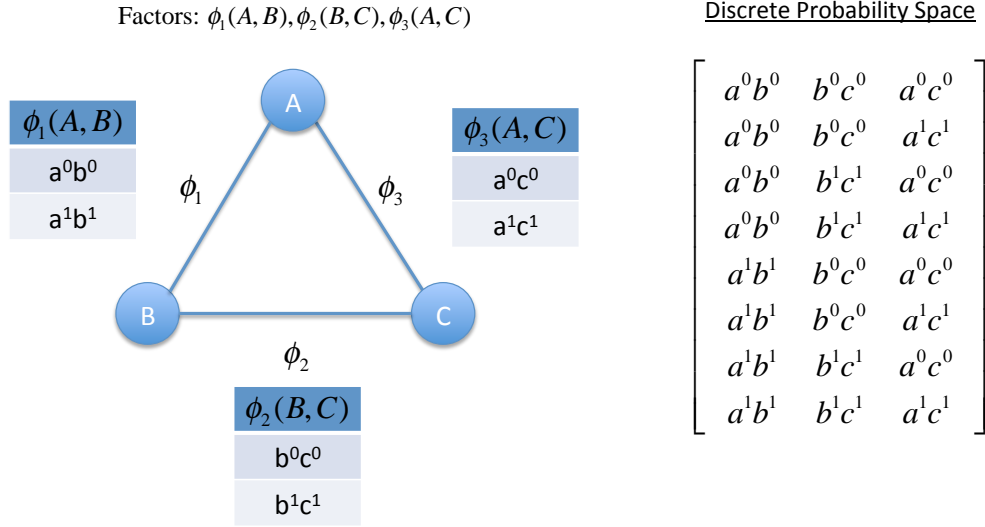


Figure 10.1: On the left is an example three-node MRF where each of the factors have two possibilities between the nodes. On the right is a representative discrete probability space enumerating the possibilities such that each row has an assigned numerical value.

The partition function creates havoc for the MRF since each time the relationship between variables change, it requires a recalculation of the probabilities throughout the network. For instances where the goal is to infer the optimal factor relationships between vertices through an optimization routine, this computation can be burdensome.

10.3.3 Pairwise Markov Networks

The Hammersley-Clifford theorem states that “..a positive distribution ($p(y) > 0$) satisfies the conditional independence of a (Markov Random Field) if and only if it can be represented as a product of factors, one per maxim clique” [158]. A pairwise MRF is defined by restricting parameterization to the edges of the network instead of the maximal cliques.

For the learning of the network structure of the GCM, let $X = (X_1, X_2, \dots, X_p)$ denote a random vector such that each of the elements are associated with the graph vertices. The Markov Random Network is the graph over the random vector X that is the family of distributions which factorize as $p(x) \propto \exp\{\sum_{(s,t) \in E} \phi(x_s, x_t)\}$, where each edge is represented by a factor relationship between the two nodes.

10.3.4 Ising Model

A particular pairwise MRF of interest is known as the Ising Model [159]. This was introduced by Ising in 1925 for applications in statistical mechanics. It is classically used for modeling interactions between positively and negatively charged particles. Of interest is an Ising model where $X_s \in -1, 1$ for each vertex and the edges are represented by the factor $\phi_{st}(x_s, x_t) = \theta_{st}^* x_s x_t$ for some parameter $\theta_{st}^* \in \mathbb{R}$. The probability distribution takes the form:

$$\mathbb{P}_{\theta^*(x)} = \frac{1}{Z(\theta^*)} \exp\{\sum \theta_{st}^* x_s x_t\} \quad (10.2)$$

This formulation is useful in learning network structure since the model can describe the the binary choice whether two nodes have an edge that is connected (1) or disconnected (-1). As previous described, the partition function ensures that the distribution sums to one.

10.3.5 Logistic Regression

Logistic regression is a generalization of linear regression for a binary distribution. This is accomplished in two steps: First, the Gaussian distribution is replaced with a Bernoulli distribution and second, a sigmoid function is used to insure that the estimate is either zero or one. This is also called a squashing function since it takes input from the domain of reals and outputs on the interval $[0,1]$. Typically, a separating criteria is established to determine which values are rounded to zero or one.

Logistic regression is a misnomer. It is really an example of a classification algorithm since the output of a logistic regression process is either a negative (0) or positive (1) result. An attribute of the logistic regression is that it is highly interpretable—the simple binary scaling permits easy human interpretation. It typically also has a small number of

parameters and is computationally efficient in solving for the parameters. The downside of the approach is its poor comparative performance with respect to other modeling techniques such as Support Vector Machines, Random Forests and Boosted Trees.

10.3.6 ℓ_1 Regularization

Sparse regularization is a process for the simplification of a regression model by removing non-influential parameters. One popular method for doing so is using the ℓ_1 norm, this is known as ℓ_1 regularization. In the context of a pairwise MRF, removing non-influential parameters has the effect of removing edges from the graph in order to provide a more succinct description of a model.

For example, suppose that there exists a linear regression model that produces an estimator and as a part of the estimator function there is a variable (x) that is relatively expensive and does not contribute to good predictions. It would be useful to drive the parameter associated with x to zero. With the parameter set to zero the corresponding x has no influence on the estimator and it is removed. This provides a less expensive model that is simpler and performs nearly as well.

10.3.7 ℓ_1 Regularized Logistic Regression

Now consider combining together ℓ_1 regularization with logistic regression. This will be modified slightly for the particular problem at hand but it is useful to start from the common model. Given a supervised learning problem with a set of n independent and identically distributed examples $\mathbf{x}_1, \dots, \mathbf{x}_p \in X$ with the associated classification label $y_1, \dots, y_n \in Y$. Logistic regression models the probability distribution of the classification label y given the feature vector \mathbf{x} :

$$p(y = 1|\mathbf{x}; \theta) = \sigma(\theta^T \mathbf{x}) = \frac{1}{1 + \exp(-\theta^T \mathbf{x})} = \frac{\exp(-\theta^T \mathbf{x})}{\exp(-\theta^T \mathbf{x}) + 1} \quad (10.3)$$

The sigmoid function ($\sigma(\cdot)$) is the aforementioned squashing function and it is defined by the far right equation. Under the Laplacian prior $p(\theta) = (\lambda/2)^N \exp(-\lambda \|\theta\|_1)$, (with $\lambda > 0$), the maximum a posteriori (MAP) estimate of the parameters θ is given by:

$$\hat{\theta} = \underset{\theta}{\operatorname{argmin}} \left(- \sum_i^N \log p(y^i | \mathbf{x}^i; \theta) + \lambda \|\theta\|_1 \right) \quad (10.4)$$

This is the optimization problem known as ℓ_1 regularized logistic regression. There is an equivalent form that is more convenient for implementation. It is given by:

$$\begin{aligned} \hat{\theta} = \underset{\theta}{\operatorname{argmin}} \left(- \sum_i^N \log p(y^i | \mathbf{x}^i; \theta) + \lambda \|\theta\|_1 \right) \\ \text{subject to } \|\theta\|_1 \leq C_\lambda \end{aligned} \quad (10.5)$$

The choice of C_λ can be made so that the optimizations are equivalent. The addition of the ℓ_1 constraint complicates the optimization problem. Essentially, this is due to the difficulty in calculating the differential of a non-continuous function. This results in a constrained optimization problem with double the number of parameters learned compared to the constraints [82]. The next section briefly introduces structural learning in graphical models.

10.4 Structural Learning in Graphical Models

Structural learning in graphical models is a discovery process where one is interested in learning the interconnections (the edges) between the graph nodes. There are normally two reasons for learning the graph structure. The first is to create a new model for probabilistic reasoning. The second application involves building the model to better understand an underlying structure of a phenomena.

In our case, we are interested in learning the structure of an undirected network through the correlation between random fields among the lattice cells of the GCM. Notice that insertion of the word undirected differentiates structural learning techniques between Bayesian and Markov Random Fields (MRF). This is significant since with Bayesian networks, it is possible to calculate local normalization within each conditional probability density. As mentioned previously, for the MRF, the partitioned function must be calculated each time there is a reassignment of the parameters.

There are several techniques for structural learning. One general methodology looks to assign a numerical value to a particular hypothesis of an optimal graph structure. A score-based approach then evaluates different hypotheses or graph models as an optimization problem. The score may take into consideration the sparseness or simplicity of the graph as well as how good the model fits the data. As one could imagine, enumerating all possibilities is computationally difficult.

10.5 GCM Inference

With the background material complete, the methodology for learning network structure within the GCM is now presented. In the chapter, “Defining The Map”, the GCM was defined within a survey space Ω as a lattice of cells where the dimension of the cell corresponds to map resolution. From each cell, the goal is to build an LCM (equivalently a random field, normally in \mathbb{R}^2) based on the signal strength measurements. The random field is simplified through a representation of a discrete number of random functions such that each represents a bearing window and the elements in the vector correspond to a set of ranges from the center of the cell.

The GCM can be naturally envisioned as a graph, where each lattice cell is a network node and the edges represent the relationship between the nodes. Figure 10.2 gives an example of the network relative to Monterey Harbor, Monterey CA. Here the survey space has been partitioned into a lattice and we are interested in the correlation between neighboring cells in the GCM.

During an initial survey, as signal measurements are received, each cell uses the SCKF for producing mean and variance across several bearing windows that correspond to stochastic estimates of RSS as a function of distance. These vectors approximate random functions for each of the bearing windows. Of interest is the correlation or conditional dependency between the random functions. The most straightforward analysis involves estimating correlation between equivalent bearing windows, but it is also possible to determine correlation between random functions of dissimilar bearing windows.

With respect to graphical models, the nodes are representing random functions. They represent the event that a bearing window random function is correlated to a bearing window to a neighboring cell. Note that in this framework it is possible to evaluate the conditional dependency between random functions within the same lattice cell. By doing this it is pos-



Figure 10.2: The above graphic displays a grid overlay that represents a map resolution of a GCM.

sible to determine if the random field is isotropic. The lattice cell would be isotropic if the vertices representing each bearing window are all connected by edges.

Consider a p -dimensional discrete random variable $X^{(t)} = (X_1^{(t)}, X_2^{(t)}, \dots, X_p^{(t)})$ where the distribution of $X^{(t)}$ is determined by a time-varying pairwise Markov Random Field. The goal is to estimate graph structure from a sample of n data points $\{x^{(i)} = (x_1^{(i)}, x_2^{(i)}, \dots, x_p^{(i)})\}_{i=1}^n$.

The raw sample data starts as signal measurements, but this isn't convenient for modeling the binary choice as to whether there is evidence (or not) for an edge between vertices. Instead a simple validation test is used for determining the criteria for evidence of an edge—whether the signal measurement validates in the model of interest. A measurement (z) is received. Using the position of the source node it is associated with a particular bearing window of a lattice cell. The measurement is used by the SCKF to update the random function. Next the measurement is checked to see if it validates within the neighbors RSS model. The validation function takes the measurement (with range r and bearing ψ from the center) and compares it with the current stochastic estimate of the neighbor node. It validates $V(z, \psi, r) = 1$ if it is within two standard deviations of the mean. Mathematically:

$$V(z, \psi, r) = \begin{cases} 1 & \text{if } 2\sigma - Z_{\Omega_u} \geq z \leq 2\sigma + Z_{\Omega_u} \\ 0 & \text{otherwise} \end{cases}$$

Note that validation is dependent on the status of the SCKF for the particular vertex. The SCKF starts with an a priori assumption with regards to the mean and covariance. In the last chapter, the a priori assumption was a Matérn function with parameters (16,20,16) and a uniform variance of $\sigma^2 = 16$. This is a fairly general starting assumption for the model and it is likely that measurements will initially validate in neighboring models. As more measurements become available, the SCKF mean and variance will improve and this tighter variance bound may restrict the network structure. The impact of the process has an effect of encouraging a fully-connected GCM network to start and gradually restricting the network structure as more information becomes available. This is analogous to assuming an isotropic model and moving to an anisotropic model (if warranted) as data becomes available.

The result of the model validation process limits entries of the data sample to

$$\mathbf{X}^{(t)} = (X_1^{(t)}, X_2^{(t)}, \dots, X_p^{(t)}) \in \mathcal{X}^p \quad (10.6)$$

and $\mathcal{X}^p = \{-1, 1\}$. It is now possible to model the validated signal data at any time t is a binary, pairwise, Ising model, MRF. Given by:

$$\mathbb{P}_{\theta}(X^{(t)}) := \frac{1}{Z(\theta^{(t)})} \exp\left(\sum \theta_{uv}^{(t)} X_u^{(t)} X_v^{(t)}\right) \quad (10.7)$$

First, $\theta_{uv}^{(t)}$ is the parameter indicating the strength of the edge connecting u and v . The terms to the right of the summation multiplies the θ edge strength and the assignment (either $\{-1, 1\}$) to the vertices X_u and X_v .

A key point of difference from prior work is the modification of the classification problem to include the SCKF filter as an temporal validation process. Normally a complete data set $D := \{x^{t_1}, x^{t_2}, \dots, x^{t_n}\}$ is used for developing the sequence of temporal networks. Instead in

this approach the dataset can be viewed as a queue such that the latest feature vector x^{t_n+1} is placed immediately in front of the prior feature vector x^{t_n} . In other words, instead of a stagnate dataset the network structure is continually using new data for potentially modifying the graph edges. This methodology is conducive to a near real-time implementation onboard vehicles.

In similar fashion to [154], a weighted non-negative kernel function is used to limit the impact of older data. This is a temporal kernel function that is directly analogous to the spatial covariance function discussed in detail in part one. It also serves a second function to reduce the computation. Any data that is old enough such that it would receive a zero kernel weighting is disregarded. The Gaussian Radial Basis Function kernel is defined as follows:

$$w^{(t)}(t_i) := K_{h_n}(t - t_i) / \sum_{i=1}^n K_{h_n}(t - t_i) \quad (10.8)$$

$$K_{h_n}(t) = \exp(-t^2/h_n) \quad (10.9)$$

10.5.1 Estimating the time-varying network

The key to this methodology is the use of a neighborhood-based ℓ_1 regularized logistic regression. Instead of using a score-based or constraint-based approach, logistic regression is run on each vertex to determine the local connected neighborhood - relative to a single vertex. Once this is completed over all vertices, the result is the learned network structure. An important point is that unlike most MRF approaches, there is no requirement for calculating the partition function and this significantly reduces the computational burden.

With $r \in V$, define a neighborhood set $\mathbb{N}^{(t)}(r) := \{s \in V | (r, s) \in E\}$ as the set of vertices that share an edge with r . Let $\theta_r^* := \{\theta_{ru}^*, u \in V \setminus r\}$ signify the parameters associated with the edges that form the neighborhood around vertex r . For the calculation of the vector θ_r^* the conditional distribution of X_r given the other neighborhood vertices $X_V = \{X_t | t \in V \setminus \{r\}\}$ is given by a form of logistic regression (and similar to the right-hand side of equation 2):

$$\mathbb{P}_{\theta_v^{(t)}}(X_r^{(t)}, X_v^{(t)}) = \frac{\exp(2X_r^{(t)} \langle \theta_v^{(t)}, X_v^{(t)} \rangle)}{\exp(2X_r^{(t)} \langle \theta_v^{(t)}, X_v^{(t)} \rangle) + 1} \quad (10.10)$$

With this setup and a dataset of time samples, the goal is the estimate of the neighborhood network parameters $\theta_v^{(t)}$ through the following optimization problem:

$$\hat{\theta}_v^{(t)} = \underset{\theta}{\operatorname{argmin}} \left(- \sum_i^N \log p(y^i | \mathbf{x}^i; \theta) + \lambda \|\theta\|_1 \right) \quad (10.11)$$

Notice this is just a modification of Equation 10.4 that takes into consideration the kernel weighting function $w^{(t)}(t_i)$. This equation also has a form convenient for encoding that is commiserate with Equation 10.5. The optimization problem is a ℓ_1 regularized logistic regression with kernel re-weighting. A popular technique for solving the optimization is the Limited-memory Broyden Fletcher Goldfarb Shanno (L-BFGS) algorithm.

10.5.2 GCM Structural Learning Inference Example

Figures 10.3a and 10.3b, are plots of two mean random fields. The plots represent RSS models of the signal environment at two neighboring grid cell locations. The left-hand side model has a strong attenuation of the RSS signal in the north/south direction and the right-hand side has the strong attenuation of the signal in the east/west direction. Notice, though, the models are strongly correlated in orthogonal directions. In both cases the RSS values have normally distributed added noise with $\sigma^2 = 6.25$.

Figures 10.4(a-h) are the respective estimates of the random functions in the cardinal directions (N,E,S,W) and include 200 plotted raw measurements. Table 10.1 shows columns 1–20 and Table 10.2 show columns 181–200 of the design matrix dataset (8 x 200) where each column that represents whether the RSS measurement validated in the neighboring nodes model. The columns represent a time sequence with the right-hand side column corresponding to the latest received measurement.

For example, column one was a RSS measurement that validated in only one of the random

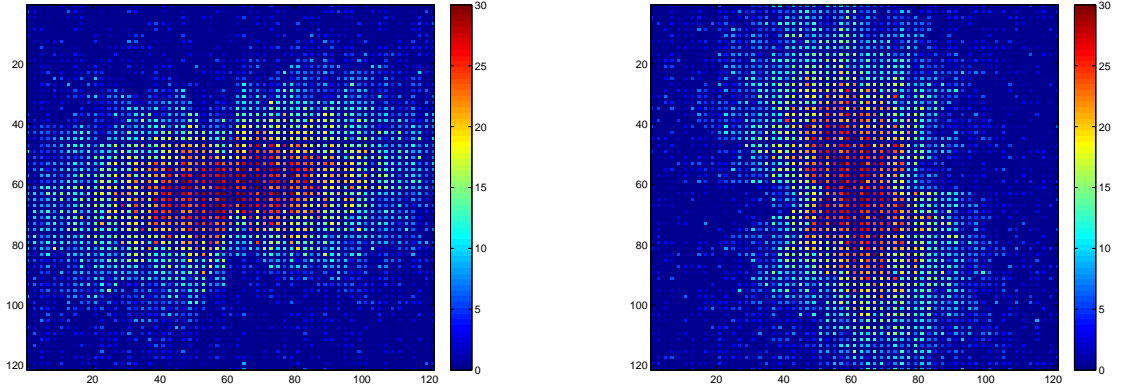


Figure 10.3: The plots represent simulated RSS acoustic modem measurements. The measurements are taken relative to a centrally located static buoy as a source node. The measurement intensity values range from 1 to 30dB using the MATLAB colormap hot.

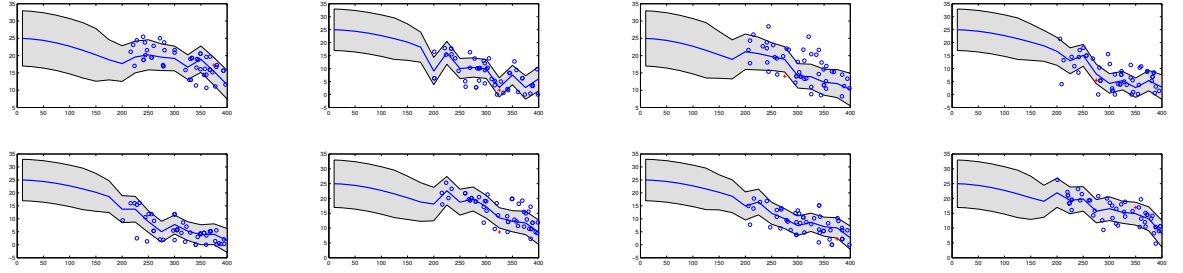


Figure 10.4: The plot shows the random functions for LCM_1 and LCM_2 given in Figure 10.3. The top row are the LCM_1 random functions for north, east, south, west and the bottom row are the same for LCM_2 .

functions while for column two, the next RSS measurement validated in all of the random functions. As time progresses, the random functions have greater accuracy and less uncertainty and no longer validate the RSS measurements as frequently. This is consistent with the shape of the models seen in Figure 10.3. The east and west random functions from LCM_1 are conditionally dependent with the north and south random functions from LCM_2 . The north and south random functions from LCM_1 are conditionally independent from the east and west random functions of LCM_2 .

Figures 10.5 and 10.6 show the accuracy and precision of the time evolving network. Even

0	1	0	1	1	1	0	1	0	1	1	1	1	1	1	0	1	1	1
0	1	0	1	1	1	0	1	0	1	1	1	0	1	1	1	0	1	1
0	1	0	1	1	1	0	1	0	1	1	1	0	1	1	1	1	1	1
0	1	0	1	1	1	0	1	0	1	1	1	0	1	1	1	0	1	1
0	1	0	1	1	1	0	1	0	1	0	1	0	1	1	1	0	1	1
1	1	0	1	1	1	0	1	1	1	1	1	1	1	1	1	1	1	1
1	1	0	1	1	1	0	1	1	1	1	1	1	1	1	1	1	1	1
0	1	0	1	1	1	0	1	0	1	1	1	1	1	1	1	0	1	0
0	1	1	1	1	1	1	1	0	1	1	1	1	1	1	1	0	1	1

Table 10.1: First 20 columns of the design matrix

0	0	1	1	0	1	0	1	0	1	0	1	0	0	1	0	1	0	1
1	1	0	1	0	1	0	1	0	0	0	1	0	1	0	1	0	1	0
0	1	1	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0
1	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	0
1	1	0	1	0	1	0	1	0	0	0	1	0	1	0	1	0	1	0
1	0	1	1	1	1	0	1	1	1	0	1	1	1	0	1	1	1	0
1	1	0	1	0	1	0	1	0	1	0	1	0	1	0	1	1	1	0
0	0	1	1	1	1	1	1	0	1	1	1	0	1	1	1	0	1	1

Table 10.2: Last 20 columns of the design matrix

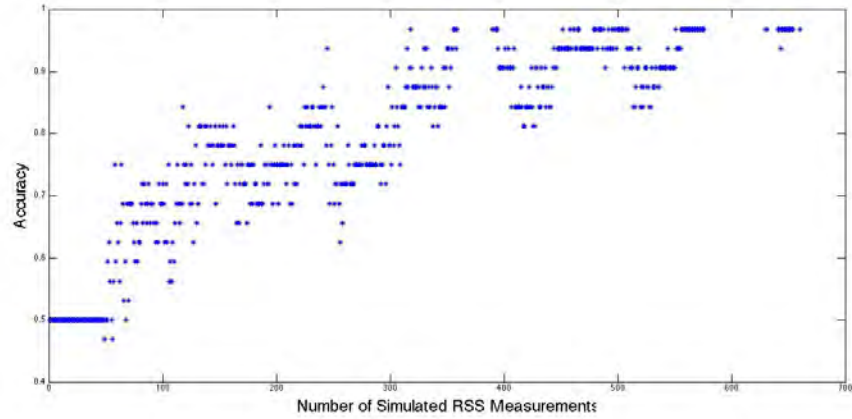


Figure 10.5: Accuracy results of the combined SCKF ℓ_1 Regularized Logistic Regression approach with the simulated example.

though there is a possibility that measurements will validate in other models, over time, the trend is for the technique to learn the proper network structure. This is shown by the upward trend in the accuracy and precision plots.

In summary, this chapter has presented a novel methodology for determining the conditional dependency between neighboring grid cells in a Global Connectivity Map. It uses an ℓ_1 regularized logistic regression together with the Spatial Communications Kriging Filter (SCKF) to temporally discover network structure. It relies on RSS measurements that are

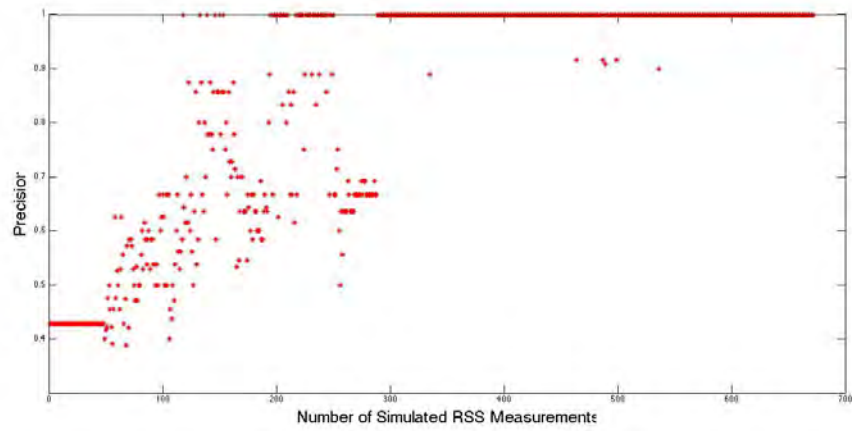


Figure 10.6: Precision results of the combined SCKF ℓ_1 Regularized Logistic Regression approach with the simulated example.

validated using the random functions that are estimated by the SCKF. The methodology is conducive to near real-time operations onboard unmanned systems.

CHAPTER 11:

GCM Experimental Results

The chapter starts with the building of a Global Connectivity Map (GCM) from under-sea acoustic modem Received Signal Strength (RSS) data collected in Monterey Harbor, Monterey in March 2013. The measurement collection strategy emphasized a smaller map resolution. The model errors are evaluated using the K-folds cross-validation techniques presented in part one. A second GCM is built from the RSS measurements collected in August 2012. This GCM reflects a collection strategy emphasizing fewer LCMs. Error analysis is again performed. The sampling strategies between the two data sets reflect different assumptions about the nature of the acoustic environment. The two models are compared and a table is presented that summarizes appropriate sampling strategies for different environments.

11.1 GCM Results (March 2013)

RSS measurements were collected in Monterey Harbor, Monterey, CA on February 28 and March 1, 2013. Over this two-day period, the following process was repeated a total of 20 times. First, a gateway buoy was placed at a designated location in the harbor. This buoy acted as the source node. A surface boat navigated around the harbor. It acted as the destination node. It traveled through the harbor with the acoustic modem deployed over the side of the boat in approximately 8-foot water depth. The micro-modem mini-packet was transmitted between the gateway buoy and the boat's acoustic modem every ten seconds. The packet summarizes the signal statistics and includes RSS, SNR In, SNR Out and noise levels. The weather was kind. There was little wind (especially in the harbor) and the surface conditions throughout the two days were calm. The tide data is given in Figure 11.1.

Figure 11.3 gives the locations of the static buoys for the two days. On the first day, there were 13 distinct sites, on the second day there were eight. Tables 11.5 and 11.6 give the Lat/Lon position of the static buoys, the number of measurements at each site and the start time. Figure 11.4 shows the combined plots of the static buoy locations. Note that repeated sites are annotated in orange and the collection on two consecutive days at the same locations provides the ability to do some initial temporal analysis.

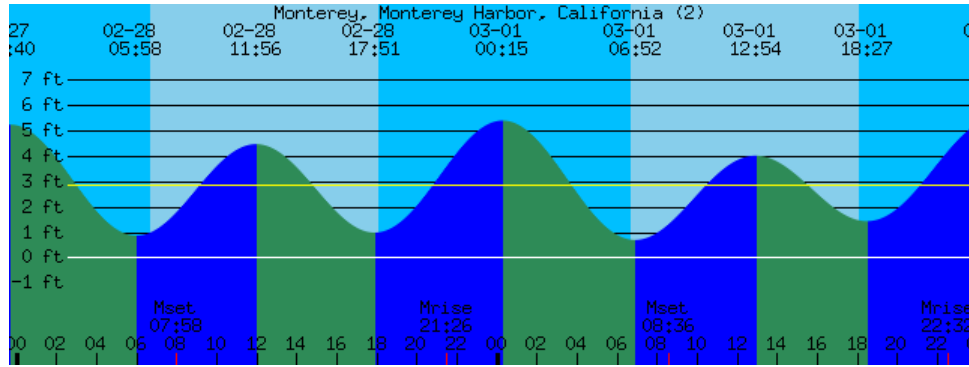


Figure 11.1: Tide chart for February 28, 2013



Figure 11.2: First day positions for the static gateway acoustic buoy in Monterey harbor.

The goal of the data collection was to create a GCM with a greater number of partitions and thereby a higher resolution GCM. This meant that fewer measurements were collected at each cell when compared with the August 2012 dataset. A deliberate choice was made to collect RSS measurements in nominal north/south, east/west directions. In other words, to the greatest degree possible, a crossing pattern. This was done as best as possible despite the numerous navigational hazards. Figure 11.7 shows an overlay of the measurements taken from the boat when the gateway buoy was located at S_{18} .

Figure 11.9 are the plots of the location and strength of the raw Received Signal Strength (RSS) measurements relative to several of the buoy locations (S_{34} , S_{35} , S_{25} , S_{45} , S_{44} , S_{43} , S_{55} , S_{54}).



Figure 11.3: Second day positions for the static acoustic gateway buoy in Monterey harbor

The GCM is a tessellation of LCMs. As such, it is necessary to select a method for determining the partition of the region taking into consideration the ability of the environment to support acoustic communication. One popular approach for dividing the configuration space is to use a Voronoi partition [160]. Given a distance metric $d : M \times M \rightarrow \mathbb{R}_{\geq 0}$, a set $S \subset M$ and n distinct points $P = \{p_1, \dots, p_n\}$ in S , the Voronoi partition of S generated by P is the collection of convex sets $V(P) = \{V_1(P), \dots, V_n(P)\} \subset \mathbf{P}(S)$ defined by, for each $i \in \{1, \dots, n\}$,

$$V_i(P) = \{q \in Q \mid \|q - p_i\| \leq \|q - p_j\|, \forall p_j \in P, i \neq j\}$$

Figure 11.10 shows the Voronoi partition of the harbor with the control points being static buoy locations throughout the collection period. In layman terms, the Voronoi partitions are formed by equal distances from the closest control points. In this case, the control points were selected that provided good coverage of the outer harbor. The partition can be considered the GCM resolution. Recall that within the boundary of each partition, it is assumed that the region can be accurately described by a single random field. This uniformity or homogeneity of the LCM is an important consideration in deriving sampling strategies and will be discussed later in greater detail.



Figure 11.4: Combined positions of the gateway acoustic buoy where orange annotates locations where data sets were collected on both days.

Recall that the GCM is designed to show 1). An estimate of overall communication effectiveness, 2). An estimate of overall communication uncertainty, 3). A directional estimate of mean and variance associated with each LCM. Figure 11.11 shows the GCM that combines the together communication effectiveness with the directional mean and variance.

The colors in each of the Voronoi partitions reflects a scalar value that is the overall estimate to communicate. The depicted scale on the right margin is calculated relative to an ideal acoustic modem environment based on an open ocean environment. It is assumed to be an isotropic model with a maximum range of 1000m and the one-dimensional view is shown in Figure 11.14. The Riemann sum of the signal model is calculated out to 300m. The color of each Voronoi partition is calculated relative to this ideal model.

Emanating from the buoy positions are a series of vectors. These represent the strength of the RSS model with respect to a bearing (0° , 45° , 90° , 135° , 180° , 225° , 270° , 315°). It

Name	Buoy Location	Number of measurements	Start Time
S_{34}	36.607352, -121.891018	417	2013 0228 072150
S_{35}	36.607346, -121.890142	107	2013 0228 080538
S_{35}	36.607346, -121.890142	333	2013 02283 082638
S_{25}	36.608043, -121.890099	408	2013 0228 091243
S_{45}	36.606609, -121.890128	363	2013 0228 100434
S_{44}	36.606616, -121.891027	352	2013 0228 104929
S_{43}	36.60664, -121.891939	630	2013 0228 113221
S_{55}	36.605906, -121.890999	335	2013 02283 125324
S_{54}	36.605899, -121.891968	167	2013 0228 142641
S_{53}	36.605901, -121.892804	314	2013 0228 144933
S_{52}	36.605205, -121.891018	290	2013 0228 153416
S_{62}	36.605207, -121.892794	267	2013 0228 160425
S_{64}	36.605205, -121.891018	189	2013 0228 164225
S_{65}	36.605197, -121.890109	154	2013 0228 171548
S_{33}	36.607351, -121.891937	134	2013 0228 174946

Figure 11.5: *Buoy Positions, the number of measurements taken and the start time for the RSS data collection in Monterey Harbor, Monterey CA February 28, 2013*

provides a depiction of the directional signal strength of the LCMs. Mathematically, the magnitude of the vector is calculated using a line integral from the estimate of the random field in that particular direction.

There are several interesting aspects to the GCM. As one might expect, the best areas for communication generally coincide with center of the harbor channel and the deepest most open regions of the harbor. This is reflected in two ways on the GCM. First, by the regions annotated in orange to reddish colors and second, by the longer signal strength vectors. The worst communication region was in the southeast of the harbor. This was somewhat unexpected. The region had relatively deep water with few communication/navigational hazards (such as surface ships). Finally, the region in the southwest had a surprisingly good ability to communicate given that it had shallow water and was relatively confined.

11.1.1 GCM Accuracy

Optimal sampling for more traditional Kriging applications such as mining, requires the use of a hexagonal or square grid pattern [14]. It also depends on the variability of the data. For Communications Kriging, in an isotropic environment, incremental sampling along a

Name	Buoy Location	Number of measurements	Start Time
S_{12}	36.608873, -121.89364	101	2013 0301 072152
S_{13}	36.608799, -121.892838	58	2013 0301 074345
S_{13}	36.608799, -121.892838	268	2013 0301 075810
S_{13}	36.608799, -121.892838	275	2013 0301 083556
S_{25}	36.608043, -121.890099	468	2013 0301 091308
S_{44}	36.606616, -121.891027	345	2013 0301 104930
S_{44}	36.606616, -121.891027	382	2013 0301 104930
S_{21}	36.608012, -121.893792	406	2013 0301 113337
S_{32}	36.607344, -121.892812	352	2013 0301 123007
S_{33}	36.607351, -121.891937	228	2013 0301 131141

Figure 11.6: *Buoy Positions, the number of measurements and the start time for the RSS data collection in Monterey Harbor, Monterey CA March 1, 2013*

single line of bearing from the source signal to a maximal designated radial distance would be sufficient for building an accurate Kriging model over the complete circular region. For anisotropic environments, adequate sampling would be dictated by the variability of the angular data.

The collected data in March 2013 reflects the objective of finer map resolution at the possible expense of the LCM accuracy. From a practical standpoint, it wasn't possible to collect the data in an optimal sampling pattern for each LCM over the short time period. As it was, it took two days to complete the data collection and the assumption that the acoustic environment remained the same over this timeframe is possibly unrealistic.

The GCM model accuracy can be evaluated using the K-folds cross-validation process. For each of the LCM “sub-models”, a K-folds process was run on each of the datasets. Figure 11.13 shows the aggregated result, that is, calculating the Kriging error for each of the LCM models and combining the results. The best fitting normal pdf has a mean ($\mu = 0.1866$) and a variance ($\sigma^2 = 3.27$) with a total of 6620 error calculations.

The GCM shows to be an accurate model but there is one note of caution—as can be seen by the Figure 11.13, there were a greater number of small Kriging errors than would expected with a normal distribution. I believe this is a limitation in the experimental design, measurements were clustered along lines particular routes such that when data was removed

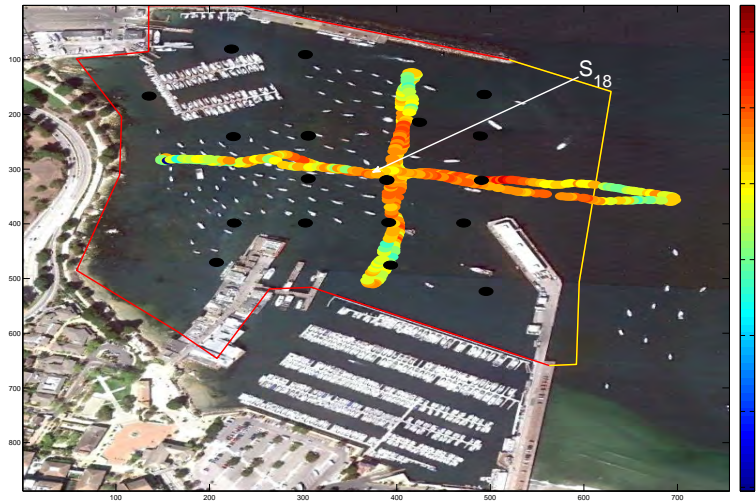


Figure 11.7: RSS Measurements collected relative to the S_{18} gateway buoy position. In black are the designated positions of the acoustic modem gateway buoy.

in the K-folds process there was still plenty of nearby points for the Kriging estimator.

11.2 GCM Results (July 2012)

The July 2012 dataset provides an opportunity for comparative analysis. The GCM resulting from the July 2012 datasets is given in Figure 11.18. It has a larger map resolution less since there were fewer Voronoi partitions, but notice that the overall ability to communicate generally conforms to the GCM in Figure 11.11. With the July 2012 GCM, the fewer partitions were offset by a greater number of measurements relative to each site. The individual LCMs built from this dataset were presented in the experimentation results section of part one. Table 11.15 provides a summary of the the number of the RSS measurements. In general, the center of the harbor shows again to be the best area to communicate, with the surrounding outer areas in the northwest and southeast the poorest.

LCM accuracy was presented in Chapter Seven. Figure 11.17 shows the combined results of Kriging errors across all LCMs. In comparing the Kriging errors associated with two models, each has a near zero mean Gaussian distribution but the variance of the models were significantly different. The variance of the histogram plots from the March 2013

were consistently larger. This is to be expected since there were fewer RSS measurements for each region. Additionally, the measurements were concentrated in cardinal directions rather than uniformly distributed throughout the region.

11.3 GCM Measurement Strategies

A natural question with respect to building the GCM is to determine the optimal sampling strategy of spatial measurements that balances time and energy with accuracy. A good strategy would determine trajectories for multiple vehicles that permits measurements in locations to simultaneously build the multiple random fields comprising the GCM.

There are two fundamental aspects of the GCM. The first is the anisotropic nature of the random field. This has already been introduced and is the uniformity of the RSS estimate with respect to the bearing from the source node. The second is the homogeneity within the Voronoi cell. The GCM makes a simplifying assumption that all estimates within the cell boundary are the same as the estimate with respect to the defined central node of the random field. This is the map resolution of the GCM.

When this assumption is valid, the cell is homogeneous. The source node could be located anywhere within the cell boundary and the resulting random field would be the same. The converse is an example of a non-uniform or non-homogeneous GCM cell. The movement of the central node to other positions within the LCM boundary would result in a significantly different estimate. In this case, the cell boundaries are poorly defined, but were made necessary, possibly, due to the limited number of available measurements.

Table 11.1 summarizes these general characteristics. Each entry reflects an efficient strategy for accurately building a GCM in a given environment. The top left entry is a GCM that is characterized by a relatively homogeneous, isotropic field. This is an ideal situation where it is possible to estimate the GCM with minimal measurements and the GCM cells are characterized by a larger area. The bottom right entry is the opposite situation, a difficult communications environment that is characterized by a non-homogeneous and anisotropic environment that requires denser measurements with a GCM cell resolution that must be smaller in order to create an accurate map.

The off-diagonal entries represent intermediate situations. The top right is an environment with a non-homogeneous, isotropic field which requires a sparser number of measure-

ments and smaller GCM cells. The top left environment is characterized by a homogeneous, anisotropic field which requires denser measurements and supports larger GCM cells.

The collection strategies that resulted in the July 2012 and March 2013 datasets reflect these off-diagonal table entries. In July 2012, only a few (5) fixed acoustic buoy locations were selected but the spatial measurements more completely covered the bounded survey region. This corresponds to the bottom left entry in the table and would be an appropriate data collection strategy for a homogeneous, anisotropic GCM. In March 2013, there were a significantly greater number of acoustic buoy locations (20), but each survey consisted of a more limited number measurements. This strategy corresponds to the top right entry of the table, one that is appropriate for an non-homogeneous, isotropic environment.

Table 11.1: *GCM data collection strategies*

	Homogeneous	Non-homogeneous
Isotropic	Sparse measurements Larger GCM map resolution	Sparse measurements Smaller GCM map resolution
Anisotropic	Dense measurements Larger GCM map resolution	Dense measurements Smaller GCM map resolution

11.4 Conclusion

In summary, the chapter presented experimental results of building GCMs. Two GCMs were presented each with an accuracy assessment based on the K-folds cross-validation methodology. The GCM created from the March 2013 data reflected a collection strategy that placed an emphasis on a greater map resolution. This was accomplished by increasing the number of Voronoi control points within the GCM, where each control point was the position of the stationary communications buoy. The downside of the approach was that the number of measurements for each LCM was less than ideal.

The second GCM was constructed from data collected in August 2012. It featured fewer Voronoi control points but this was offset by a greater number of measurements within each LCM. The comparison of the two GCM highlighted sampling strategies for different

environments. Given knowledge of an environment it is possible to design, in advance, more efficient strategies for building the GCM.

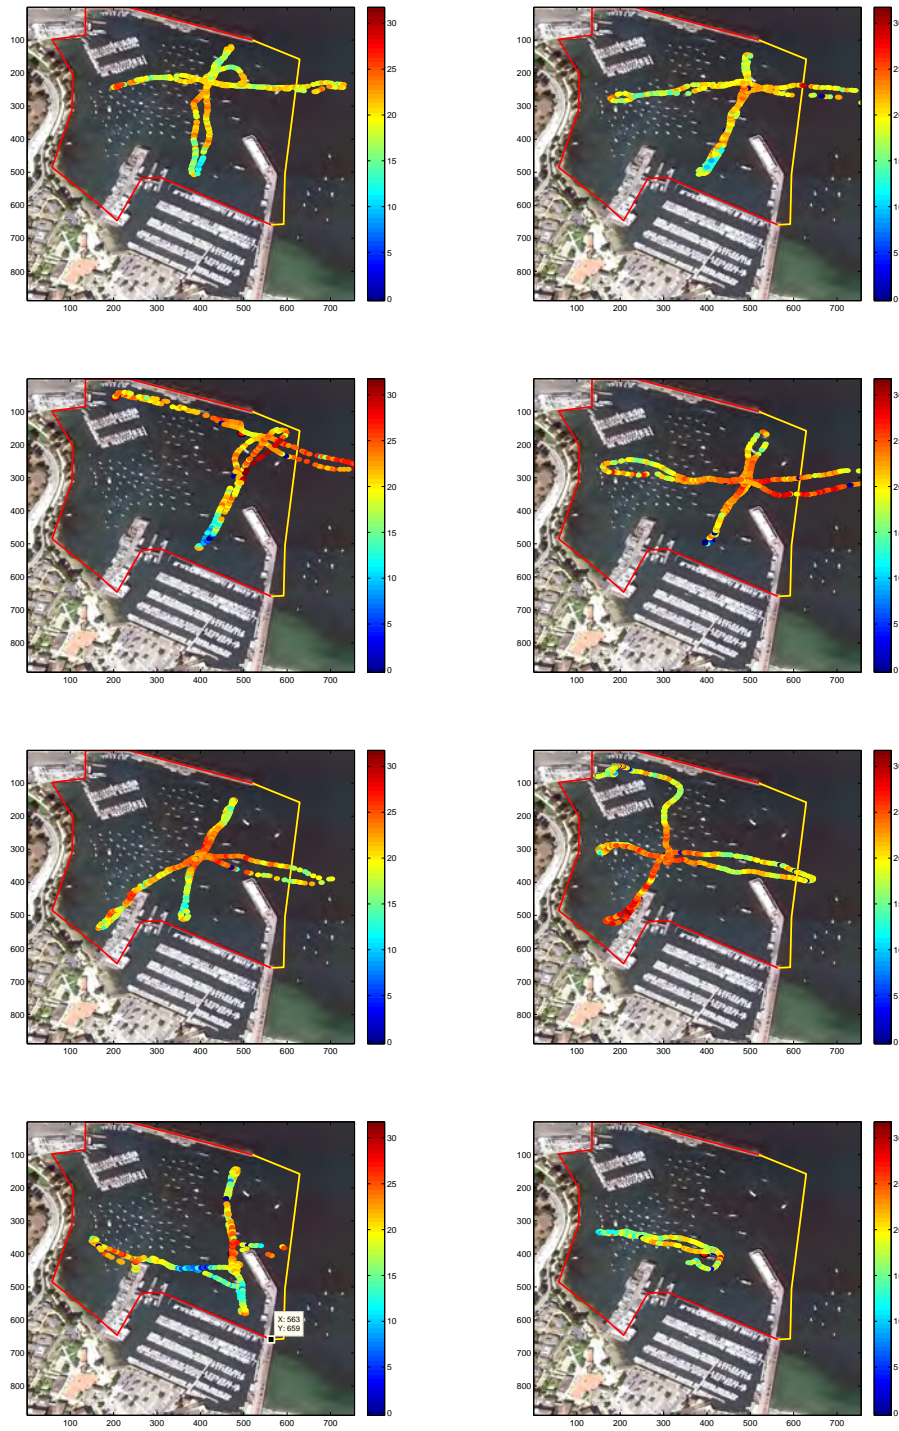


Figure 11.8: RSS measurements collected on February 28, 2013

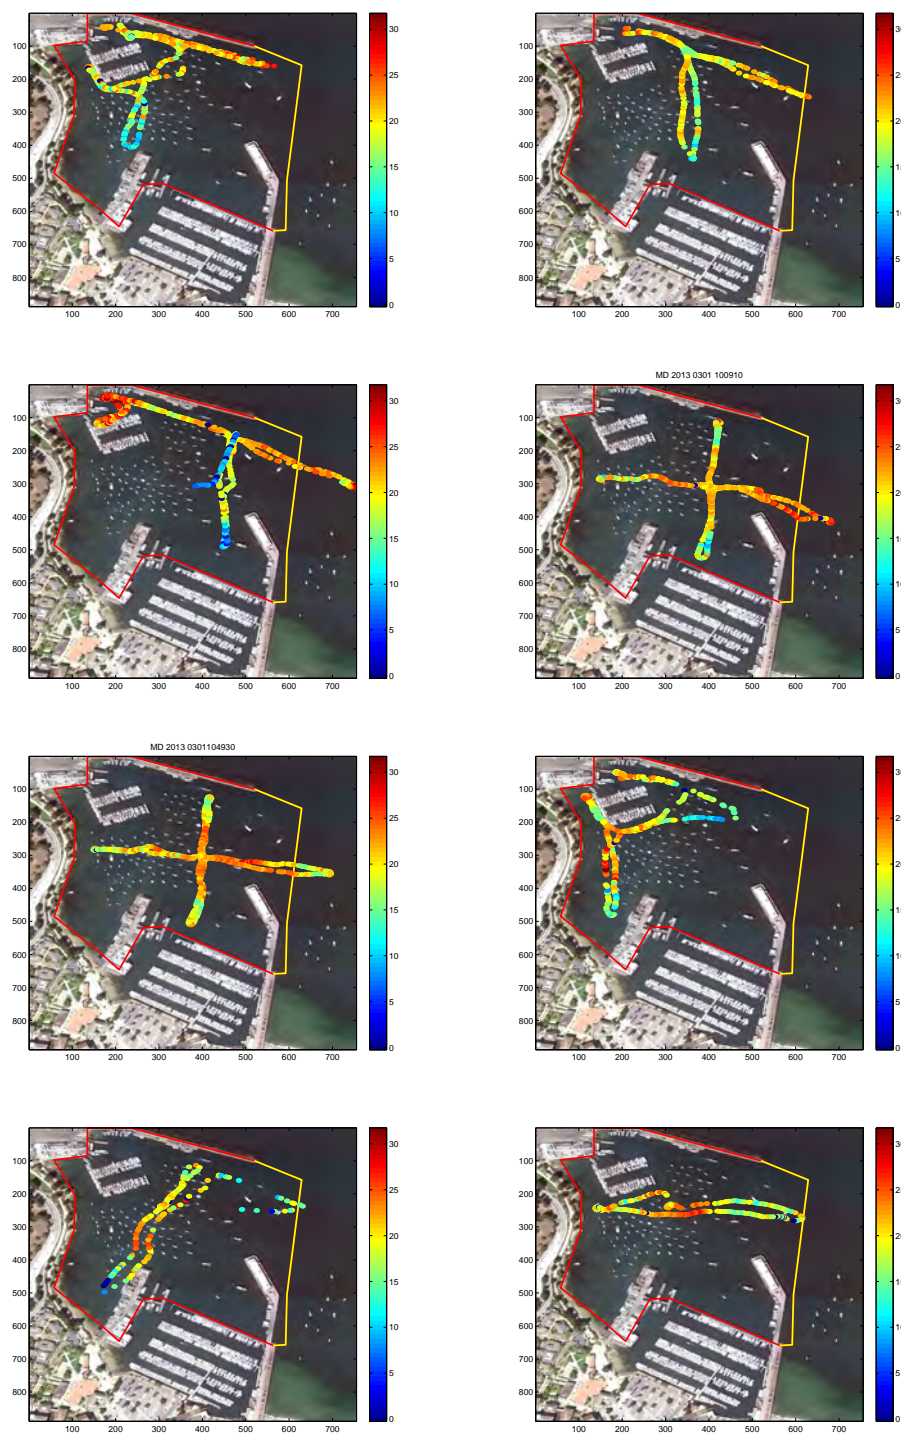


Figure 11.9: *RSS measurements collected on March 1, 2013*



Figure 11.10: A voronoi partition with the static buoy positions acting as the control points of the partition

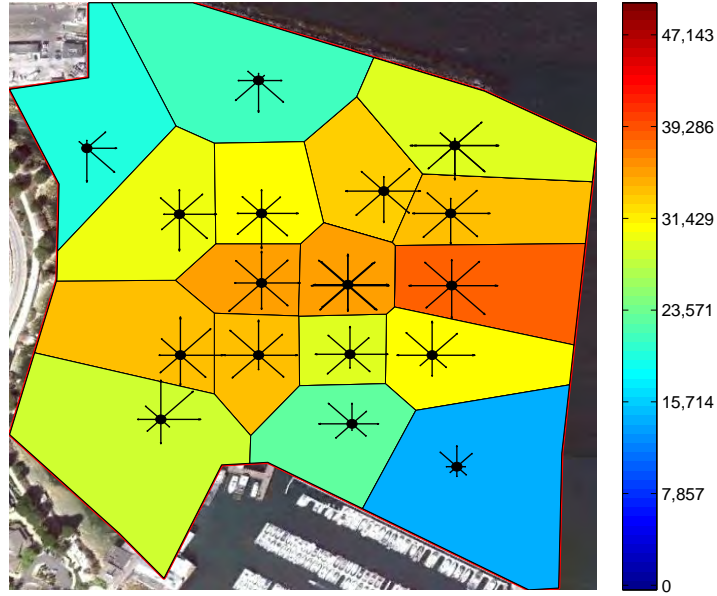


Figure 11.11: A GCM based on March 2013 RSS acoustic modem data collected in Monterey Harbor, Monterey, CA.

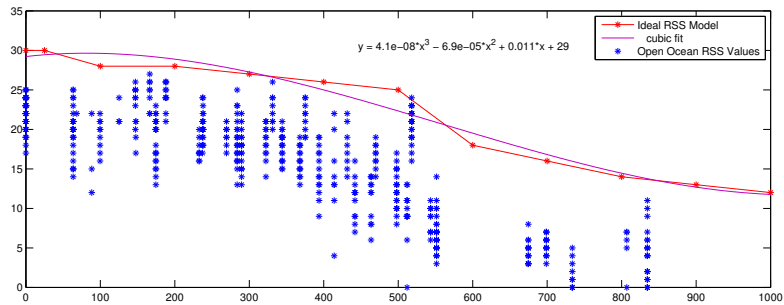


Figure 11.12: The LCM reference model. The random fields of the GCM are compared to the reference model to determine the map color scaling.

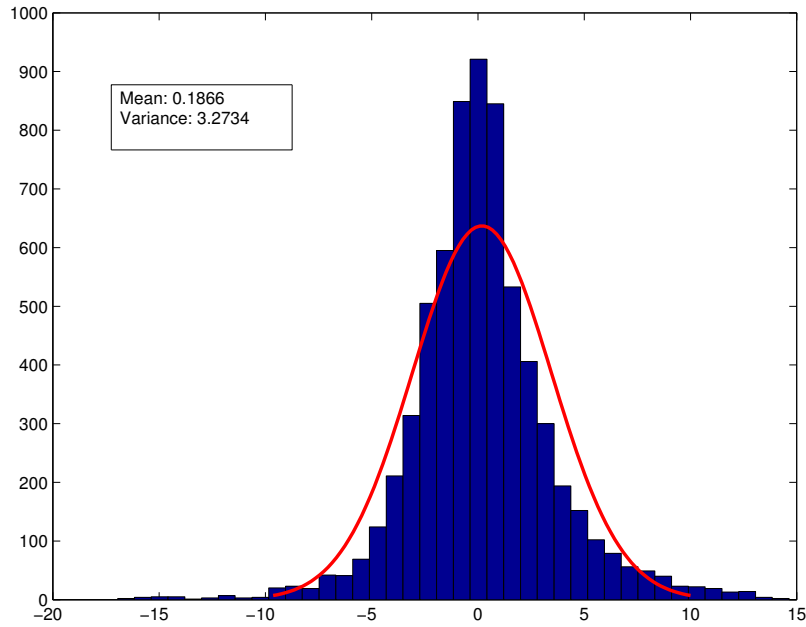


Figure 11.13: Results of the K-Fold cross-validation method. The red curve is a best fitting normal distribution with $\mu = 0.1866$ and $\sigma^2 = 3.27$.

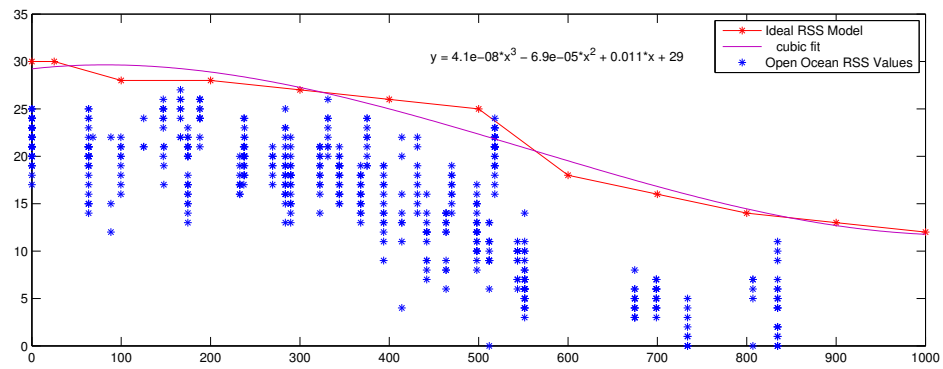


Figure 11.14: The function in magenta provides the basis for calculating the LCM reference model. It is used to determine the GCM scalar value for each partition.

Name	Buoy Location	Number of measurements	Start Time
S_1	36.607024, -121.889747	963	30 Jul 2012 18:49:01 GMT
S_2	36.606500, -121.891680	752	30 Jul 2012 20:33:08 GMT
S_3	36.607755, -121.893768	525	30 Jul 2012 22:12:11 GMT
S_4	36.608948, -121.892566	500	30 Jul 2012 17:49:59 GMT
S_5	36.608917, -121.892535	589	30 Jul 2012 16:54:21 GMT

Figure 11.15: RSS Measurements (July 2012)

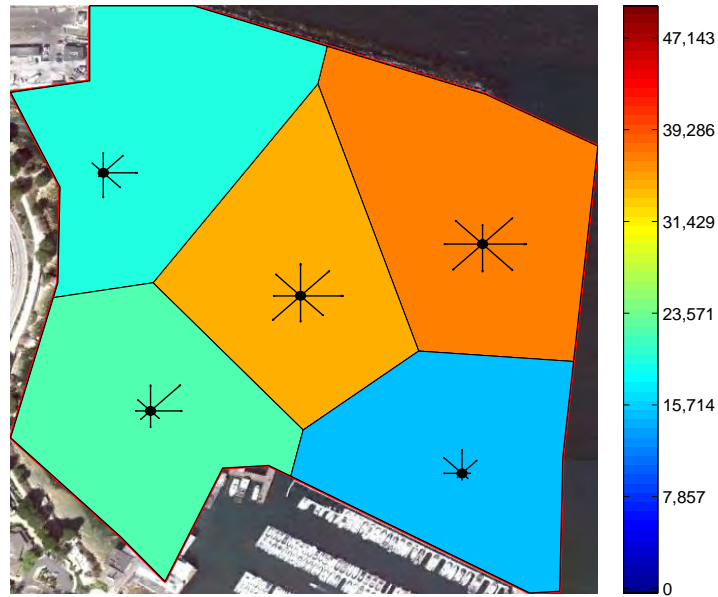


Figure 11.16: The July 2012 GCM.

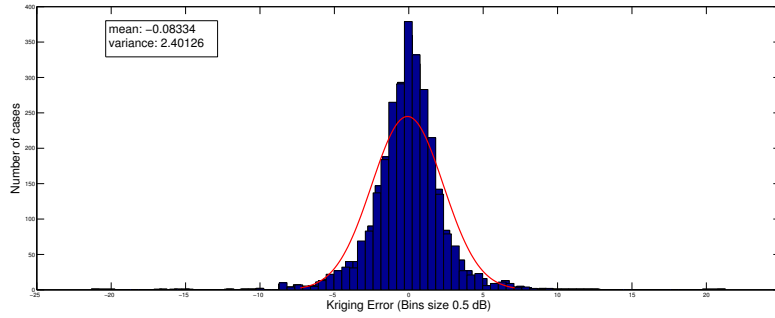


Figure 11.17: Results of the K-Fold cross-validation method for the July 2012 GCM. The red curve is a best fitting normal distribution with $\mu = -0.08$ and $\sigma^2 = 2.40$.

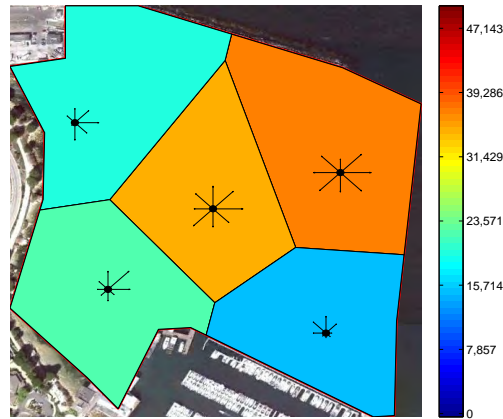


Figure 11.18: July 2012 GCM. The colors correspond to general assessments of signal strength and the vectors centered on the central cell node correspond to directional signal strength.

THIS PAGE INTENTIONALLY LEFT BLANK

CHAPTER 12:

Conclusions and Recommendations

In a variety of domains, there is a requirement for wireless signal strength estimation. State of the art approaches consist of an attenuation or physics-based model that frequently does not adequately represent the true signal propagation. This dissertation presented a rigorous mathematical framework that supports efficient construction of connectivity maps with minimal measurements. It emphasized a data-driven methodology. The goals were three-fold:

1. Provide greater resolution or understanding with regards to signal propagation as a function of range and bearing.
2. Provide an overarching framework to support stochastic wireless signal strength estimation between either static or mobile nodes, especially in dynamic and cluttered environments.
3. Design the framework to support near real-time operations.

The two-part presentation introduced channel estimation from a local and global perspective. These definitions are consistent with concepts of local and global reference frames frequently used in robotics and control literature. The first part of the research presented the Local Connectivity Map (LCM). The LCM provides a channel estimates relative to the pose of an agent. The second part presented the Global Connectivity Map (GCM). It provides an estimate of signal strength between two or more positions within a region—regardless of position.

This chapter presents a review the approach and briefly summarizing the results. It also includes recommendations for future work. It starts with the LCM.

12.1 LCM

It is initially assumed that there are a set of signal strength measurements over an area. The sampling is incomplete and for this reason it is necessary to infer the signal strength at all locations (to a defined map resolution) based on nearby measurements. The construction of the LCM leveraged the theory of Random Fields and a technique in optimal spatial es-

timization known as Kriging. Kriging is an interpolation technique where the first step is determining the spatial dependence of data. This is known as variogram analysis. Typical implementations of Kriging use a Euclidean distance measurement between random variables. A novel contribution of the research extended Kriging, through the recognition that, with a source-based signal, it is possible to leverage the nature of the signal for improving the stochastic estimate.

The first step was the conversion of the data into a range and bearing relative to a source node—in other words, conversion of the data into a polar distance metric. This permits the creation of two distinct semi-variograms—a range and radial semi-variogram. This permits a range and radial estimate that can be combined to create the final mean and variance at a POI. The combination of techniques for a source-based signal are collectively called Communications Kriging.

These techniques were demonstrated using undersea acoustic modem RSS measurements taken in Monterey Harbor, Monterey, CA. A static acoustic modem buoy was positioned at designated sites and a boat with a second acoustic modem navigated the harbor to collect RSS measurements. This was done twice in August 2011 and July 2012. Using cross-validation techniques, the LCMs built with the Communication Kriging techniques showed good performance—the mean errors were low, the variances were consistent with the measurement errors of the semi-variogram (the nugget).

12.2 GCM

The second part of the dissertation focused on the Global Connectivity Map (GCM). It provides a signal strength estimate between any two points in a bounded area. While the LCM was relative to the pose of an agent in the network, the GCM is independent of vehicle position. For this reason it is called a *global* map.

The GCM was defined as a lattice of LCMs. A key to the definition of the GCM was a relaxation of the quasi-stationarity assumption. With respect to the LCM, the quasi-stationarity assumption defines a neighborhood of random *variables* with equivalent mean values. For the GCM, within each cell border, the quasi-stationarity assumption for random *fields* is assumed to hold, such that the region can be summarized by a single random field.

The goal of the GCM is simultaneous estimation of a lattice of LCMs. Only in this way,

is it possible to estimate RSS between source and destination between any location within the GCM. To build the GCM, it is necessary to build estimates of the LCMs in near real-time. This implies the use of a recursive filtering solution for LCM estimation. The Spatial Communication Kalman Filter (SCKF) was introduced for this purpose. It combines aspects of optimal spatial estimation with recursive filtering to provide an estimate of RSS as a function of range (from the center of the lattice cell). The result is a mean and variance estimate of a vector of ranges that is relevant to a select bearing window.

The GCM uses the information provided by the filters to summarize the ability to communicate. It does in two ways, by calculating a Riemann sum of the mean over the several bearing windows to produce an estimate of the ability to communicate and calculates a normalized Riemann sum of the variance to produce an uncertainty estimate.

The lattice cells are arbitrarily placed over the survey region. The size of the survey regions dictate the map resolution. In the GCM results section, the lattice cells were constructed as Voronoi cells. It is of interest to determine the similarity between regions. If it were determined that two lattice cells were similar, this information could be used to improve the boundary definition of lattice cells. The implication is that by combining cells into an aggregate cell, the number of measurements needed for the GCM could be reduced. Second, a longer term goal is the ability to determine environmental phenomenon that impacts the ability to communicate.

To accomplish these objectives it is necessary to establish a framework for determining the similarity between cells. This was introduced in Chapter Ten—Refining the GCM. The chapter uses a neighborhood-based ℓ_1 regularized logistic regression for learning the time-varying network structure of a pairwise Markov Random Field. The novelty of the approach is the adaptation of the supervised learning technique for near real-time estimation by the use of the SCKF as a validation function for calculating valid vector entries into the ℓ_1 regularized logistic regression design matrix.

This involves two aspects. The first is determining a feature vector that can be used as input for estimating graph structure. The ℓ_1 regularized logistic regression typically uses an Ising model where graph nodes are encoded either by $\{1, -1\}$. For the GCM this is accomplished by using the current random function estimate by the SCKF. If a RSS measurement validates in a neighboring graph node, it is encoded with a 1 and -1 otherwise.

It was shown that this simple validation process has a couple of implications. The ability to learn the graph model is based on the current filter estimate provided by the SCKF. In the beginning, with a non-informative prior with wide variance, measurements will tend to easily validate in the model. As the model becomes more accurate, the resulting impact of the graph structure will be to limit the edges between nodes for a more informative graph model. Results were presented using two simulated RSS LCMs. It demonstrated the ability of the technique to properly learn the correct network structure.

12.3 Model Comparison

It is difficult to directly compare the data-driven approach to the physics-based modeling approach. This is due to the limitations of physics-based modeling to handle complex communication environments. The emphasis of the research was the development of a technique that is appropriate for cluttered and dynamic environments. This was illustrated for undersea acoustic communications by the selection of Monterey Harbor as the data collection site.

A state of the art physics-based model for undersea acoustic communications uses a ray-tracing algorithm based on parameters that can include water temperature, bathymetry and bottom-type. While this may be appropriate for the open ocean, it is not considered an accurate model for a congested harbor. It does not usually take into consideration the multi-path and fading that occurs as a result of the structures and moving objects that reside in the harbor.

It is believed that a data-driven modeling approach could be combined together with a signal attenuation and physics-based modeling approach to bring about better LCMs (and GCMs). This is especially true in the underwater domain with unmanned vehicles. AUVs typically have onboard sensors to evaluate the environment. The signal attenuation model can take into consideration transducer design and the impact of the mounting the transducer on the AUV. Recent research combined a signal attenuation and physics-based modeling approach for undersea acoustic communication [161] and [162]. Future work could combine the data driven model with this approach.

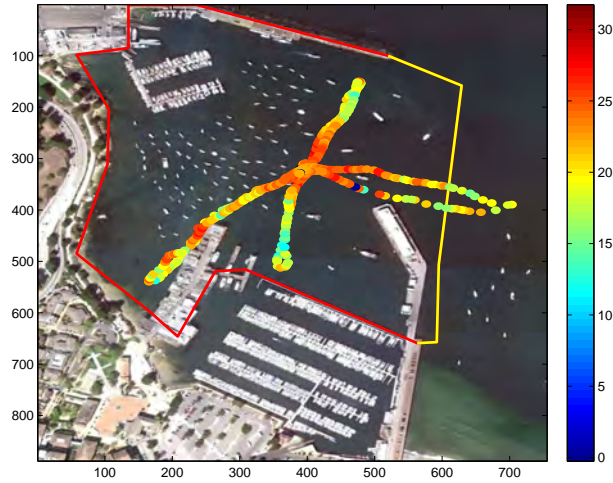


Figure 12.1: Source node S_{44} . Data collected on February 28, 2013

12.4 Temporal Model Considerations

In Chapter Ten, a network structure learning algorithm was presented for learning dependencies between neighboring random fields. A strong dependence between two random fields could reduce the number of partitions in the GCM. Comparisons were calculated through a validation process driven by the SCKF random field estimates. As measurements became available over time, they were used to improve the GCM accuracy. This provided a rudimentary temporal model. But it was only good for the time that the measurements were collected. A natural question to ask is how long are the maps good for. Due to limitations in equipment, acoustic datasets were collected over the course of an entire day. It may be reasonable to assume that the map remained consistent for the entire collection period.

Figure 12.1 and 12.2 and show examples of datasets that were collected at three different times. Each time the acoustic modem gateway buoy was at the same location. While the measurement locations were not coincident, the degradation of the RSS over distance is similar. A caveat is that the weather over these two days was remarkably good and consistent. There was little wind or wind-generated waves.

This example brings out several points: First, under benign conditions, the RSS measurements can be used to produce LCMs/GCMs that can stay consistent over a longer time

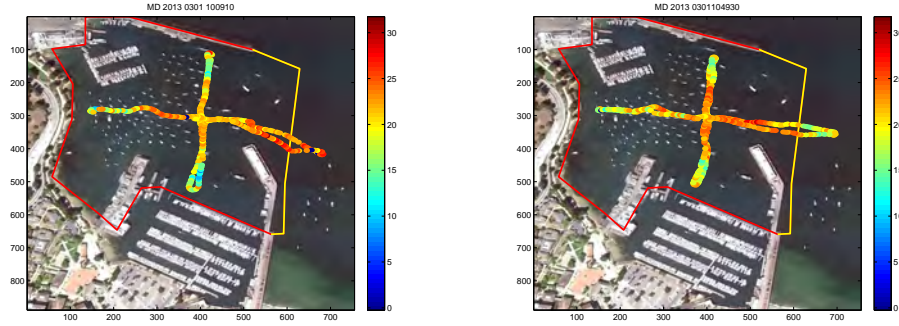


Figure 12.2: Source node S_{44} . (a) RSS Data collected at 1009 PST on March 01, 2013, (b) RSS Data collected at 1049 PST on March 01, 2013

frame (in this case about 24 hours). This implies that measurements can be potentially used over the entire time period for a single realization of the random field(s). Second, the consistency of the LCM/GCM is dependent on the environment. For the acoustic undersea signal strength measurements, it is suspected that, if the surface winds increased and created wind-driven waves, it would impact the LCM/GCM and thereby impact the consistency of the temporal estimate.

Third, consideration of the temporal aspect of the measurements is an important consideration for future work. Difficult questions would include how to partition the measurements over time periods that reflect changing environmental considerations. In the above example, one possibility would be to have a state within the SCKF filter that estimated surface winds. The combination of the LCM/GCM estimation together with environmental modeling could yield an explanatory hypothesis that could be used for partitioning temporal data.

12.5 GCM Techniques

In Chapter Ten, the goal was to better define the boundaries of the GCM. The technique addressed the ability to simplify the map by potentially joining areas with similar signal strength. This is an aggregation strategy. It did not address a technique to better determine border regions through a “refinement” strategy—a further partitioning of a single cell border based on measurements that validates a hypothesis of multiple random fields within the cell border.

These techniques are two of a set of processes can be collectively viewed to handle issues in GCM accuracy, speed of construction and robustness. With respect to accuracy, there are at least two salient aspects. The first is the accuracy of the LCMs. This was discussed in Chapter Seven, in terms of the Kriging error associated with the model using K-folds cross-validation. The second is the accurate placement of the partition borders. One way to state the problem is that the assumption with respect to the quasi-stationarity within the cell partition was incorrect. Given that there were enough measurements, it might be possible to construct experimentation using significance testing.

Another interesting potential approach is the recognition that the central node locations for the GCM cells are model parameters. Viewed in this way, it may be possible to use traditional techniques for determining optimal values for the central cell positions. For example, setting up an optimization of the parameters using a negative log likelihood function.

Additionally, one could use a past GCM as a prior in constructing an updated version of the model in the same way that the Matérn function was used as a prior for the SCKF. Viewing a past GCM as a prior, given it is accurate, will permit a new GCM to convergence more quickly with potentially less measurements. Collectively these approaches are considered as future research items.

12.6 Additional Future Research

There are a number of areas for additional future research. These are presented with respect to the organization of the dissertation.

12.6.1 LCM

1. Fundamental to the success of the Kriging technique is the variogram. It determines the accuracy of the connectivity map. The techniques presented centered on processing data after the data had been collected. A potentially better way to do this would be to conduct a survey that takes the dual objective of structural analysis on data sampling into consideration. The result would be an information-theoretic framework for path-planning that simultaneously reduced uncertainty about the semi-variogram(s) and reduced map error through intelligence selection of the sampling points.
2. More analysis is warranted with regards to functions for communication semi-variograms. Common parametric methods didn't fit the semi-variogram data particu-

larly well and the shape of the semi-variogram was highly dependent on the spacing between lag distances and the choice of a maximum range. Non-parametric methods provide better elasticity but in question is whether the data truly reflects the underlying relationship between data at different lag distances.

3. The building of the LCM focused on the RSS as the single signal measurement to assess the ability to communicate. This is a simplistic approach. A better implementation would take several variables into consideration. Using a series of signal statistic variables that included RSS, SNR, packet error rate and others might produce a better metric to assess the probability of successfully passing a message.
4. Now that it has been demonstrated how the LCM estimation process can be run in near real-time, the next step would be to use the results of the LCM to provide useful information for control feedback for collaborative navigation. Additionally the LCM can be used to determine dynamic rate scaling where the modem to set a rate for data transfer based on the LCM estimate.

12.6.2 GCM

1. Building the GCM is a difficult process due to the number of measurements necessary to estimate multiple random fields. Similar to the intelligent autonomy techniques described the LCM section—the multi-vehicle path planning problem for the quick construction of the GCM is an interesting problem. The collective paths of the vehicles should be constructed such as to simultaneously expand and contract communications ranges from the nominal grid cell centers.
2. The SCKF was demonstrated emphasizing range estimation. The SCKF could be also used for radial estimation. Additionally, the SCKF could include estimation of the radial and range covariance parameters. The combination would provide a real-time technique for estimation of the random field.
3. The method presented for learning graph structure in the GCM made a simplifying assumption with regards to the validation process. If the measurement validated in another neighboring node the feature vector was assigned a one and zero otherwise. The measurement was taken at a single range. This validation process is rather narrow. It could be made more robust with a different validation technique.
4. As previously mentioned, the combination of the presented techniques provides a general methodology for detecting an explanatory environmental phenomena. A collaborative system of agents could collectively use the wireless communications

ability as a kind of second-order sensing system.

12.7 Conclusion

The presented work has been an opportunity to investigate spatial estimation techniques as it pertains to communication. It has been interesting to compare and contrast these techniques with (the possibly) more traditional temporal estimation processes that are commonly used for navigational filters. It is the belief of the author that spatial estimation will have an increasing role within robotics applications. The main reason is transitioning of robotics from navigation-centric research to work-oriented objectives.

In that migration, of fundamental importance is the robot's internal representation of the world. Spatial estimation techniques are a natural way to use sparse measurements for constructing an internal environmental model. Furthermore, the combination of spatial and temporal techniques are a natural way to model cluttered and dynamic environments. This will be especially critical for robotic systems to interoperate amongst people in challenging environments.

THIS PAGE INTENTIONALLY LEFT BLANK

REFERENCES

- [1] Y. Mostofi and P. Sen, “Compressed mapping of communication signal strength,” in *Military Communications Conference, 2008. MILCOM 2008*. IEEE, 2008, pp. 1–7.
- [2] J. Hespanha, P. Naghshtabrizi, and Y. Xu, “A survey of recent results in networked control systems,” *Proceedings of the IEEE*, vol. 95, no. 1, pp. 138–162, 2007.
- [3] T. Yang, “Networked control system: A brief survey,” in *IEEE Proceedings Control Theory and Applications*, vol. 153, no. 4. IET, 2006, pp. 403–412.
- [4] R. Gupta and M. Chow, “Networked control system: Overview and research trends,” *Industrial Electronics, IEEE Transactions on*, vol. 57, no. 7, pp. 2527–2535, 2010.
- [5] Q. Lindsey, D. Mellinger, and V. Kumar, “Construction with quadrotor teams,” *Autonomous Robots*, vol. 33, no. 3, pp. 323–336, 2012.
- [6] J. Fax and R. Murray, “Graph laplacians and stabilization of vehicle formations,” *15th IFAC Congress, Barcelona, Spain*, 2002.
- [7] R. Olfati-Saber and R. M. Murray, “Consensus problems in networks of agents with switching topology and time-delays,” *IEEE Transactions on Automatic Control*, vol. 49, no. 9, pp. 1520–1533, 2004.
- [8] M. Porfiri and D. Stilwell, “Consensus seeking over random weighted directed graphs,” *IEEE Transactions on Automatic Control*, vol. 52, no. 9, pp. 1767–1773, 2007.
- [9] E. Vanmarcke, *Random Fields: Analysis and Synthesis*. MIT Press, Cambridge, MA, 2010.
- [10] R. J. Adler and J. E. Taylor, *Random Fields and Geometry*. Springer, New York, NY, 2007, vol. 115.
- [11] N. Cressie, *Statistics for Spatial Data (Wiley Series in Probability and Statistics)*. Wiley-Interscience, New York, NY, 1993.
- [12] J. Journel, A.G. Huijbregts, *Mining Geostatistics*. Academic Press, London, UK, 1976.
- [13] B. Ripley, *Spatial Statistics*. Wiley-Interscience, New York, NY, 2005, vol. 575.
- [14] C. Jean-Paul and D. Pierre, *Geostatistics: Modeling Spatial Uncertainty*. John Wiley & Sons Inc., New York, 1999.

- [15] M. Stein, *Interpolation of Spatial Data: Some Theory for Kriging*. Springer Verlag, New York, NY, 1999.
- [16] D. Krige, “A statistical approach to some basic mine valuation problems on the witwatersrand,” *Journal of the Chemical, Metallurgical and Mining Society of South Africa*, vol. 52, pp. 119–139, 1953.
- [17] P. Curran, “The semivariogram in remote sensing: an introduction,” *Remote Sensing of Environment*, vol. 24, no. 3, pp. 493–507, 1988.
- [18] A. Balaguer, L. Ruiz, T. Hermosilla, and J. Recio, “Definition of a comprehensive set of texture semivariogram features and their evaluation for object-oriented image classification,” *Computers & Geosciences*, vol. 36, no. 2, pp. 231–240, 2010.
- [19] M. Stein, “Nonstationary spatial covariance functions,” Center for Integrating Statistical and Environmental Science, University of Chicago, IL, USA, Tech. Rep., 2005.
- [20] K. Haskard, “An anisotropic matern spatial covariance model: Reml estimation and properties.” Ph.D. dissertation, University of Adelaide, School of Agriculture, Food and Wine, 2007.
- [21] P. Hall, N. I. Fisher, and B. Hoffmann, “On the nonparametric estimation of covariance functions,” *The Annals of Statistics*, vol. 22, pp. 2115–2134, 1994.
- [22] K. Yu, J. Mateu, and E. Porcu, “A kernel-based method for nonparametric estimation of variograms,” *Statistica Neerlandica*, vol. 61, no. 2, pp. 173–197, 2007.
- [23] C. Huang, T. Hsing, and N. Cressie, “Nonparametric estimation of the variogram and its spectrum,” *Biometrika*, vol. 98, no. 4, pp. 775–789, 2011.
- [24] M. Umer, L. Kulik, and E. Tanin, “Kriging for localized spatial interpolation in sensor networks,” in *Scientific and Statistical Database Management*. Springer, 2008, pp. 525–532.
- [25] M. Umer, L. Kulik, and E. Tanin, “Spatial interpolation in wireless sensor networks: localized algorithms for variogram modeling and kriging,” *Geoinformatica*, vol. 14, no. 1, pp. 101–134, 2010.
- [26] G. Lu and D. Wong, “An adaptive inverse-distance weighting spatial interpolation technique,” *Computers & Geosciences*, vol. 34, no. 9, pp. 1044–1055, 2008.
- [27] G. Lin and L. Chen, “A spatial interpolation method based on radial basis function networks incorporating a semivariogram model,” *Journal of Hydrology*, vol. 288, no. 3, pp. 288–298, 2004.

- [28] D. Dardari, A. Conti, U. Ferner, A. Giorgetti, and M. Win, "Ranging with ultrawide bandwidth signals in multipath environments," *Proceedings of the IEEE*, vol. 97, no. 2, pp. 404–426, 2009.
- [29] Y. Shen and M. Win, "Fundamental limits of wideband localization part i: a general framework," *Information Theory, IEEE Transactions on*, vol. 56, no. 10, pp. 4956–4980, 2010.
- [30] H. Shen, Y. Wymeersch and M. Win, "Fundamental limits of wideband localization part ii: cooperative networks," *Information Theory, IEEE Transactions on*, vol. 56, no. 10, pp. 4981–5000, 2010.
- [31] M. Gholami, "Positioning algorithms for wireless sensor networks," Ph.D. dissertation, Chalmers University of Technology, Sweden, 2011.
- [32] H. Larkin, "Wireless signal strength topology maps in mobile adhoc networks, embedded and ubiquitous computing," in *International Conference EUC2004*, 2004, pp. 538–547.
- [33] D. Cottingham, R. Harle, and A. Hopper, "Constructing accurate, space-efficient, wireless coverage maps for vehicular contexts," in *Proceedings of the 4th Annual International Conference on Wireless Internet*. ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2008, p. 44.
- [34] D. Cottingham, "Vehicular wireless communication," Computer Laboratory, University of Cambridge, Cambridge, UK, Tech. Rep., 2009.
- [35] T. Pogel and L. Wolf, "Analysis of operational 3g network characteristics for adaptive vehicular connectivity maps," in *Wireless Communications and Networking Conference Workshops (WCNCW), 2012*. IEEE, 2012, pp. 355–359.
- [36] J. Biswas and M. Veloso, "Wifi localization and navigation for autonomous indoor mobile robots," in *2010 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2010, pp. 4379–4384.
- [37] K. Yedavalli and B. Krishnamachari, "Sequence-based localization in wireless sensor networks," *IEEE Transactions on Mobile Computing*, vol. 7, no. 1, pp. 81–94, 2008.
- [38] A. Bo, Z. Zhang-dui, Z. Gang, and L. Jian-ping, "Novel statistical processing methods for wireless field strength prediction," *IEEE Transactions on Consumer Electronics*, vol. 55, no. 4, pp. 1805–1809, 2009.
- [39] M. Erol-Kantarci, H. Mouftah, and S. Oktug, "A survey of architectures and localization techniques for underwater acoustic sensor networks," *Communications Surveys & Tutorials, IEEE*, vol. 13, no. 3, pp. 487–502, 2011.

- [40] G. Han, J. Jiang, L. Shu, Y. Xu, and F. Wang, "Localization algorithms of underwater wireless sensor networks: A survey," *Sensors*, vol. 12, no. 2, pp. 2026–2061, 2012.
- [41] H. Durrant-Whyte and T. Bailey, "Simultaneous localisation and mapping (SLAM): Part I: the essential algorithms," *Robotics and Automation Magazine*, 2006.
- [42] T. Bailey and H. Durrant-Whyte, "Simultaneous localisation and mapping (SLAM): Part II: state of the art," *Robotics and Automation Magazine*, 2006.
- [43] B. Ferris, D. Hähnel, and D. Fox, "Gaussian processes for signal strength-based location estimation," in *In Proc. of Robotics Science and Systems*. Citeseer, 2006.
- [44] B. Ferris, D. Fox, and N. Lawrence, "Wifi-slam using gaussian process latent variable models," in *Proceedings of the 20th International Joint Conference on Artificial Intelligence*, 2007, pp. 2480–2485.
- [45] J.-S. Gutmann, G. Brisson, E. Eade, P. Fong, and M. Munich, "Vector field slam," in *IEEE International Conference on Robotics and Automation (ICRA)*, 2010. IEEE, 2010, pp. 236–242.
- [46] E. Candès and M. Wakin, "An introduction to compressive sampling," *Signal Processing Magazine, IEEE*, vol. 25, no. 2, pp. 21–30, 2008.
- [47] Y. Mostofi, M. Malmirchegini, and A. Ghaffarkhah, "Estimation of communication signal strength in robotic networks," in *2010 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2010, pp. 1946–1951.
- [48] Y. Mostofi, "Compressive cooperative sensing and mapping in mobile networks," *IEEE Transactions on Mobile Computing*, vol. 10, no. 12, pp. 1769–1784, 2011.
- [49] J. Haupt, R. Castro, and R. Nowak, "Distilled sensing: Adaptive sampling for sparse detection and estimation," *IEEE Transactions on Information Theory*, vol. 57, no. 9, pp. 6222–6235, 2011.
- [50] E. Arias-Castro, E. Candes, and M. Davenport, "On the fundamental limits of adaptive sensing," *arXiv preprint arXiv:1111.4646*, 2011.
- [51] D. Needell and R. Vershynin, "Signal recovery from incomplete and inaccurate measurements via regularized orthogonal matching pursuit," *IEEE Journal of Selected Topics in Signal Processing*, vol. 4, no. 2, pp. 310–316, 2010.
- [52] N. Patwari, J. Ash, S. Kyperountas, A. Hero III, R. Moses, and N. Correal, "Locating the nodes: cooperative localization in wireless sensor networks," *Signal Processing Magazine, IEEE*, vol. 22, no. 4, pp. 54–69, 2005.

- [53] J. Strom and E. Olson, "Multi-sensor attenuation estimation (matte): Signal-strength prediction for teams of robots," in *2012 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*. IEEE, 2012, pp. 4730–4736.
- [54] J. Fink and V. Kumar, "Online methods for radio signal mapping with mobile robots," in *2010 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2010, pp. 1940–1945.
- [55] S. Zickler and M. Veloso, "Rss-based relative localization and tethering for moving robots in unknown environments," in *2010 IEEE International Conference on Robotics and Automation (ICRA)*. IEEE, 2010, pp. 5466–5471.
- [56] H. Omre, "Bayesian kriging - merging observations and qualified guesses in kriging," *Mathematical Geology*, vol. 19, no. 1, pp. 25–39, 1987.
- [57] P. Kitanidis, "Parameter uncertainty in estimation of spatial functions: Bayesian analysis," *Water resources research*, vol. 22, no. 4, pp. 499–507, 1986.
- [58] H. Omre and K. Halvorsen, "The bayesian bridge between simple and universal kriging," *Mathematical Geology*, vol. 21, no. 7, pp. 767–786, 1989.
- [59] J. Hartigan, "Linear bayesian methods," *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 446–454, 1969.
- [60] R. Kulkarni, "Bayesian kriging in geotechnical problems," *Geostatistics for Natural Resources Characterization Part*, vol. 2, 1984.
- [61] J. Pilz and G. Spöck, "Why do we need and how should we implement bayesian kriging methods," *Stochastic Environmental Research and Risk Assessment*, vol. 22, no. 5, pp. 621–632, 2008.
- [62] G. E. Box and G. C. Tiao, *Bayesian Inference in Statistical Analysis*. John Wiley & Sons, 2011, vol. 40.
- [63] M. Handcock and M. Stein, "A bayesian analysis of kriging," *Technometrics*, vol. 35, no. 4, pp. 403–410, 1993.
- [64] M. Gaudard, M. Karson, E. Linder, and D. Sinha, "Bayesian spatial prediction," *Environmental and Ecological Statistics*, vol. 6, no. 2, pp. 147–171, 1999.
- [65] J. Berger, V. De Oliveira, and B. Sansó, "Objective bayesian analysis of spatially correlated data," *Journal of the American Statistical Association*, vol. 96, no. 456, pp. 1361–1374, 2001.
- [66] J. Berger, J. Bernardo, and D. Sun, "The formal definition of reference priors," *The Annals of Statistics*, vol. 37, no. 2, pp. 905–938, 2009.

- [67] R. Paulo, “Default priors for gaussian processes,” *The Annals of Statistics*, vol. 33, no. 2, pp. 556–582, 2005.
- [68] C. Rasmussen and C. Williams, *Gaussian Processes for Machine Learning*. MIT press Cambridge, MA, 2006, vol. 1.
- [69] M. Katzfuss, “Bayesian nonstationary spatial modeling for very large datasets,” *Environmetrics*, 2013.
- [70] M. Katzfuss and N. Cressie, “Bayesian hierarchical spatio-temporal smoothing for very large datasets,” *Environmetrics*, vol. 23, no. 1, pp. 94–107, 2012.
- [71] D. Koller and N. Friedman, *Probabilistic Graphical Models: Principles and Techniques*. MIT press, Cambridge, MA, 2009.
- [72] H. Loeliger, “An introduction to factor graphs,” *Signal Processing Magazine, IEEE*, vol. 21, no. 1, pp. 28–41, 2004.
- [73] F. Kschischang, B. Frey, and H. Loeliger, “Factor graphs and the sum-product algorithm,” *IEEE Transactions on Information Theory*, vol. 47, no. 2, pp. 498–519, 2001.
- [74] B. Frey, “Extending factor graphs so as to unify directed and undirected graphical models,” in *Proceedings of the Nineteenth conference on Uncertainty in Artificial Intelligence*. Morgan Kaufmann Publishers Inc., 2002, pp. 257–264.
- [75] R. Kindermann, J. L. Snell *et al.*, *Markov Random Fields and Their Applications*. American Mathematical Society Providence, RI, 1980, vol. 1.
- [76] J. Pearl, *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufmann Pub, 1988.
- [77] A. Dogandzic and B. Zhang, “Distributed estimation and detection for sensor networks using hidden markov random field models,” *IEEE Transactions on Signal Processing*, vol. 54, no. 8, pp. 3200–3215, 2006.
- [78] J. Lafferty, A. McCallum, and F. Pereira, “Conditional random fields: Probabilistic models for segmenting and labeling sequence data,” *International Conference on Machine Learning*, 2001.
- [79] C. Sutton and A. McCallum, “An introduction to conditional random fields for relational learning,” *Introduction to statistical relational learning*, vol. 93, pp. 142–146, 2007.
- [80] C. Sutton, “An introduction to conditional random fields,” *Foundations and Trends in Machine Learning*, vol. 4, no. 4, pp. 267–373, 2012.

- [81] C. Chow and C. Liu, “Approximating discrete probability distributions with dependence trees,” *Information Theory, IEEE Transactions on*, vol. 14, no. 3, pp. 462–467, 1968.
- [82] S.-I. Lee, V. Ganapathi, and D. Koller, “Efficient structure learning of markov networks using ℓ_1 -regularization,” in *Advances in neural information processing systems*, 2006, pp. 817–824.
- [83] M. Schmidt, K. Murphy, G. Fung, and R. Rosales, “Structure learning in random fields for heart motion abnormality detection,” *CVPR. IEEE Computer Society*, 2008.
- [84] M. Schmidt, “Graphical model structure learning with ℓ_1 -regularization,” Ph.D. dissertation, University of British Columbia, 2010.
- [85] D. Shahaf, A. Chechotka, and C. Guestrin, “Learning thin junction trees via graph cuts,” *Artificial Intelligence and Statistics (AISTATS)*, vol. 3, no. 3, p. 2, 2009.
- [86] F. Bullo, J. Cortés, and S. Martinez, *Distributed Control of Robotic Networks: a Mathematical Approach to Motion Coordination Algorithms*. Princeton University Press, 2009.
- [87] W. Ren and Y. Cao, *Distributed Coordination of Multi-Agent Networks: Emergent Problems, Models, and Issues*. Springer, New York, NY, 2010, vol. 83.
- [88] M. Mesbahi and M. Egerstedt, *Graph Theoretic Methods in Multiagent Networks*. Princeton University Press, 2010.
- [89] Y. Mostofi, A. Gonzalez-Ruiz, A. Gaffarkhah, and D. Li, “Characterization and modeling of wireless channels for networked robotic and control systems-a comprehensive overview,” in *IEEE/RSJ International Conference on Intelligent Robots and Systems, 2009. IROS 2009*. IEEE, 2009, pp. 4849–4854.
- [90] J. Fink, A. Ribeiro, and V. Kumar, “Robust control for mobility and wireless communication in cyber–physical systems with application to robot teams,” *Proceedings of the IEEE*, vol. 100, no. 1, pp. 164–178, 2012.
- [91] H. Wymeersch, J. Lien, and M. Z. Win, “Cooperative localization in wireless networks,” *IEEE Proceedings*, vol. 97, no. 2, pp. 427–450, 2009.
- [92] N. Trawny, S. Roumeliotis, and G. Giannakis, “Cooperative multi-robot localization under communication constraints,” in *IEEE International Conference on Robotics and Automation, 2009. ICRA’09*. IEEE, 2009, pp. 4394–4400.

- [93] X. Zhou and S. Roumeliotis, "Robot-to-robot relative pose estimation from range measurements," *IEEE Transactions on Robotics*, vol. 24, no. 6, pp. 1379–1393, 2008.
- [94] M. Zavlanos, A. Ribeiro, and G. Pappas, "Mobility & routing control in networks of robots," in *49th IEEE Conference on Decision and Control (CDC), 2010*. IEEE, 2010, pp. 7545–7550.
- [95] J. Le Ny, A. Ribeiro, and G. Pappas, "Adaptive communication-constrained deployment of unmanned vehicle systems," *IEEE Journal on Selected Areas in Communications*, vol. 30, no. 5, pp. 923–934, 2012.
- [96] J. Le Ny and G. Pappas, "Adaptive deployment of mobile robotic networks."
- [97] N. Leonard, D. Paley, R. Davis, D. Fratantoni, F. Lekien, and F. Zhang, "Coordinated control of an underwater glider fleet in an adaptive ocean sampling field experiment in monterey bay," *Journal of Field Robotics*, vol. 27, no. 6, pp. 718–740, 2010.
- [98] E. Fiorelli, N. Leonard, P. Bhatta, D. Paley, R. Bachmayer, and D. Fratantoni, "Multi-auv control and adaptive sampling in monterey bay," *Oceanic Engineering, IEEE Journal of*, vol. 31, no. 4, pp. 935–948, 2006.
- [99] D. Paley, F. Zhang, and N. Leonard, "Cooperative control for ocean sampling: The glider coordinated control system," *Control Systems Technology, IEEE Transactions on*, vol. 16, no. 4, pp. 735–744, 2008.
- [100] O. S. Erol-Kantarci M., Mouftah H., "Localization techniques for underwater acoustic sensor networks," in *IEEE Communications Magazine*, December 2010, pp. 152–158.
- [101] A. Bahr, J. J. Leonard, and M. F. Fallon, "Cooperative localization for autonomous underwater vehicles," *The International Journal of Robotics Research*, vol. 28, no. 6, pp. 714–728, 2009.
- [102] J. Cortes, "Distributed Kriged Kalman filter for spatial estimation," *IEEE Transactions on Automatic Control*, vol. 54, no. 12, pp. 2816–2827, 2009.
- [103] M. Stojanovic, "Underwater acoustic communication," *Wiley Encyclopedia of Electrical and Electronics Engineering*, 1999.
- [104] M. Stojanovic and J. Preisig, "Underwater acoustic communication channels: Propagation models and statistical characterization," *IEEE Communications Magazine*, vol. 47, no. 1, pp. 84–89, 2009.

- [105] K. J. Partan J. and L. B., “A survey of practical issues in underwater networks,” in *WUWNet06: Proceedings of the First ACM International Workshop on Underwater Networks*, 2006, pp. 17–24.
- [106] J. Heidemann, W. Ye, J. Wills, A. Syed, and Y. Li, “Research challenges and applications for underwater sensor networking,” in *Wireless Communications and Networking Conference, 2006. WCNC 2006. IEEE*, vol. 1. IEEE, 2006, pp. 228–235.
- [107] J. Heidemann, M. Stojanovic, and M. Zorzi, “Underwater sensor networks: applications, advances and challenges,” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, vol. 370, no. 1958, pp. 158–175, 2012.
- [108] P. Qarabaqi and M. Stojanovic, “Statistical modeling of a shallow water acoustic communication channel,” in *Proc. Underwater Acoustic Measurements Conference, Nafplion, Greece*, 2009.
- [109] A. Radošević, J. Proakis, and M. Stojanovic, “Statistical characterization and capacity of shallow water acoustic channels,” in *OCEANS 2009-EUROPE*. IEEE, 2009, pp. 1–8.
- [110] A. Stefanov and M. Stojanovic, “Design and performance analysis of underwater acoustic networks,” *IEEE Journal on Selected Areas in Communications*, vol. 29, no. 10, pp. 2012–2021, 2011.
- [111] L. Freitag, “The WHOI micro-modem: An acoustic communications and navigation system for multiple platforms,” *Proc. OCEANS 2005, Washington, DC*, 2005.
- [112] L. Freitag and S. Singh, “Performance of micro-modem psk signaling under variable conditions during the 2008 race and space experiments,” in *OCEANS 2009, MTS/IEEE Biloxi-Marine Technology for Our Future: Global and Local Challenges*. IEEE, 2009, pp. 1–8.
- [113] E. Gallimore, J. Partan, I. Vaughn, S. Singh, J. Shusta, and L. Freitag, “The whoi micromodem-2: A scalable system for acoustic communications and networking,” in *MTS/IEEE OCEANS 2010*. IEEE, 2010, pp. 1–7.
- [114] G. Matheron, “Principles of geostatistics,” *Economic geology*, vol. 58, no. 8, pp. 1246–1266, 1963.
- [115] A. N. Kolmogorov, “Dissipation of energy in locally isotropic turbulence,” in *Dokl. Akad. Nauk SSSR*, vol. 32, no. 1, 1941, pp. 16–18.

- [116] A. N. Kolmogorov, “The local structure of turbulence in incompressible viscous fluid for very large reynolds numbers,” in *Dokl. Akad. Nauk SSSR*, vol. 30, no. 4, 1941, pp. 299–303.
- [117] M. Armstrong and R. Jabin, “Variogram models must be positive-definite,” *Mathematical Geology*, vol. 13, no. 5, pp. 455–459, 1981.
- [118] A. Shapiro and J. Botha, “Variogram fitting with a general class of conditionally nonnegative definite functions,” *Computational Statistics & Data Analysis*, vol. 11, no. 1, pp. 87–96, 1991.
- [119] S. Bochner, *Harmonic Analysis and the Theory of Probability*. Courier Dover Publications, 1955.
- [120] M. J. Pyrcz and C. V. Deutsch, “Spectral corrected semivariogram models,” *Mathematical Geology*, vol. 38, no. 7, pp. 891–899, 2006.
- [121] M. R. Powojowski, “Isotropic spectral additive models of the covariogram,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 70, no. 4, pp. 739–753, 2008.
- [122] Y. Matsuda and Y. Yajima, “Fourier analysis of irregularly spaced data on rd,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 71, no. 1, pp. 191–217, 2009.
- [123] N. E. LeMay, “Variogram modeling and estimation,” Ph.D. dissertation, University of Colorado at Denver, 1998.
- [124] D. D. Sarma, *Geostatistics with Applications in Earth Sciences*. Springer, 2009.
- [125] R. Furrer, M. G. Genton, and D. Nychka, “Covariance tapering for interpolation of large spatial datasets,” *Journal of Computational and Graphical Statistics*, vol. 15, no. 3, 2006.
- [126] N. Cressie and G. Johannesson, “Fixed rank kriging for very large spatial data sets,” *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, vol. 70, no. 1, pp. 209–226, 2008.
- [127] R. J. Barnes and T. B. Johnson, “Positive kriging,” *Geostatistics for natural resources characterization*, no. 122, p. 231, 1984.
- [128] M. L. Stein, “Predicting integrals of random fields using observations on a lattice,” *The Annals of Statistics*, vol. 23, no. 6, pp. 1975–1990, 1995.

- [129] M. L. Stein, "Asymptotic properties of centered systematic sampling for predicting integrals of spatial processes," *The Annals of Applied Probability*, pp. 874–880, 1993.
- [130] N. Durrande, D. Ginsbourger, O. Roustant, and L. Carraro, "Additive covariance kernels for high-dimensional gaussian process modeling," *arXiv preprint arXiv:1111.6233*, 2011.
- [131] T. Hastie and R. Tibshirani, *Generalized Additive Models*. Chapman & Hall/CRC, 1990, vol. 43.
- [132] W. Jakes, *Mobile Radio Propagation*. Wiley-IEEE Press, 2009.
- [133] L. Freitag, *Micro-Modem Software Interface Guide*, Woods Hole Oceanographic Institution, 2009.
- [134] P. Bahl and V. N. Padmanabhan, "Radar: An in-building rf-based user location and tracking system," in *INFOCOM 2000. Nineteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 2. Ieee, 2000, pp. 775–784.
- [135] C. Gentile and L. Klein-Berndt, "Robust location using system dynamics and motion constraints," in *Communications, 2004 IEEE International Conference on*, vol. 3. IEEE, 2004, pp. 1360–1364.
- [136] A. M. Ladd, K. E. Bekris, A. Rudys, L. E. Kavraki, and D. S. Wallach, "Robotics-based location sensing using wireless ethernet," *Wireless Networks*, vol. 11, no. 1-2, pp. 189–204, 2005.
- [137] J. Yin, Q. Yang, and L. M. Ni, "Learning adaptive temporal radio maps for signal-strength-based location estimation," *IEEE Transactions on Mobile Computing*, vol. 7, no. 7, pp. 869–883, 2008.
- [138] A. Haeberlen, E. Flannery, A. M. Ladd, A. Rudys, D. S. Wallach, and L. E. Kavraki, "Practical robust localization over large-scale 802.11 wireless networks," in *Proceedings of the 10th annual international conference on Mobile computing and networking*. ACM, 2004, pp. 70–84.
- [139] L. M. Ni, Y. Liu, Y. C. Lau, and A. P. Patil, "Landmarc: indoor location sensing using active rfid," *Wireless networks*, vol. 10, no. 6, pp. 701–710, 2004.
- [140] Y.-C. Cheng, Y. Chawathe, A. LaMarca, and J. Krumm, "Accuracy characterization for metropolitan-scale wi-fi localization," in *Proceedings of the 3rd international conference on Mobile systems, applications, and services*. ACM, 2005, pp. 233–245.

- [141] T. Kamakaris and J. V. Nickerson, “Connectivity maps: Measurements and applications,” in *System Sciences, 2005. HICSS’05. Proceedings of the 38th Annual Hawaii International Conference on*. IEEE, 2005, pp. 307–307.
- [142] E. Bergfeldt, S. Ekelin, and J. M. Karlsson, “Real-time bandwidth measurements over mobile connections,” *European Transactions on Telecommunications*, vol. 22, no. 6, pp. 255–267, 2011.
- [143] C. Schoenfelder and S. Cambanis, “Random designs for estimating integrals of stochastic processes,” *The Annals of Statistics*, pp. 526–538, 1982.
- [144] J. Ko and D. Fox, “GP-Bayesfilters: Bayesian filtering using gaussian process prediction and observation models,” *Autonomous Robots*, vol. 27, no. 1, pp. 75–90, 2009.
- [145] J. Ko and D. Fox, “Gp-bayesfilters: Bayesian filtering using gaussian process prediction and observation models,” in *Intelligent Robots and Systems, 2008. IROS 2008. IEEE/RSJ International Conference on*. IEEE, 2008, pp. 3471–3476.
- [146] K. V. Mardia, C. Goodall, E. J. Redfern, and F. J. Alonso, “The kriged kalman filter,” *Test*, vol. 7, no. 2, pp. 217–282, 1998.
- [147] M. Hancock and J. Wallis, “An approach to statistical spatial-temporal modeling of meteorological fields (with discussion),” *Journal of the American Statistical Association*, vol. 89, pp. 368–390, 1994.
- [148] R. E. Kalman *et al.*, “A new approach to linear filtering and prediction problems,” *Journal of basic Engineering*, vol. 82, no. 1, pp. 35–45, 1960.
- [149] A. Gelb, *Applied Optimal Estimation*. MIT press, 1974.
- [150] S. Durrleman and R. Simon, “Flexible regression models with cubic splines,” *Statistics in Medicine*, vol. 8, no. 5, pp. 551–561, 1989.
- [151] P. Holme and J. Saramäki, “Temporal networks,” *Physics reports*, vol. 519, no. 3, pp. 97–125, 2012.
- [152] P. Ravikumar, M. J. Wainwright, and J. D. Lafferty, “High-dimensional ising model selection using l1-regularized logistic regression,” *The Annals of Statistics*, vol. 38, no. 3, pp. 1287–1319, 2010.
- [153] M. Kolar, L. Song, A. Ahmed, and E. P. Xing, “Estimating time-varying networks,” *The Annals of Applied Statistics*, vol. 4, no. 1, pp. 94–123, 2010.
- [154] L. Song, M. Kolar, and E. P. Xing, “Keller: estimating time-varying interactions between genes,” *Bioinformatics*, vol. 25, no. 12, pp. i128–i136, 2009.

- [155] L. Song, M. Kolar, and E. P. Xing, “Time-varying dynamic bayesian networks,” in *Advances in Neural Information Processing Systems*, 2009, pp. 1732–1740.
- [156] M. Kolar, L. Song, and E. P. Xing, “Sparsistent learning of varying-coefficient models with structural changes,” in *Advances in Neural Information Processing Systems*, 2009, pp. 1006–1014.
- [157] O. Banerjee, L. El Ghaoui, and A. d’Aspremont, “Model selection through sparse maximum likelihood estimation for multivariate gaussian or binary data,” *The Journal of Machine Learning Research*, vol. 9, pp. 485–516, 2008.
- [158] K. P. Murphy, *Machine Learning: A Probabilistic Perspective*. The MIT Press, 2012.
- [159] E. Ising, “Beitrag zur theorie des ferromagnetismus,” *Zeitschrift für Physik A Hadrons and Nuclei*, vol. 31, no. 1, pp. 253–258, 1925.
- [160] A. Okabe, B. Boots, K. Sugihara, and S. N. Chiu, *Spatial Tessellations: Concepts and Applications of Voronoi Diagrams*. Wiley. com, 2009, vol. 501.
- [161] T. Schneider and H. Schmidt, “Unified command and control for heterogeneous marine sensing networks,” *Journal of Field Robotics*, vol. 27, no. 6, pp. 876–889, 2010.
- [162] T. Schneider, “Advances in integrating autonomy with acoustic communications for intelligent networks of marine robots,” Ph.D. dissertation, Joint Program in Applied Ocean Science and Engineering at the Massachusetts Institute of Technology and the Woods Hole Oceanographic Institution, 2013.

THIS PAGE INTENTIONALLY LEFT BLANK

Initial Distribution List

1. Defense Technical Information Center
Ft. Belvoir, Virginia
2. Dudley Knox Library
Naval Postgraduate School
Monterey, California